

Fast Spectral Multiplication for Real-Time Rendering

by

C. Allen Waddle

A.Sc., General Science, Camosun College, 2005

B.Sc., Combined Physics & Biochemistry (with Distinction), University of Victoria, 2009

B.Sc., Combined Computer Science & Mathematics (with Distinction), University of Victoria, 2009

A Dissertation Submitted in Partial Fulfillment of the
Requirements for the Degree of

DOCTOR OF PHILOSOPHY

in the Department of Computer Science

© C. Allen Waddle, 2018

University of Victoria

All rights reserved. This dissertation may not be reproduced in whole or in part, by photocopying or other means, without the permission of the author.

Fast Spectral Multiplication for Real-Time Rendering

by

C. Allen Waddle

A.Sc., General Science, Camosun College, 2005

B.Sc., Combined Physics & Biochemistry (with Distinction), University of Victoria, 2009

B.Sc., Combined Computer Science & Mathematics (with Distinction), University of Victoria, 2009

Supervisory Committee

Dr. George Tzanetakis, Supervisor
(Department of Computer Science)

Dr. Alexandra Branzan Albu, Departmental Member
(Department of Computer Science)

Dr. Wu-Sheng Lu, Outside Member
(Department of Electrical and Computer Engineering)

Dr. Michael H. Brill, Additional Member
(Datacolor, New Jersey, USA)

Supervisory Committee

Dr. George Tzanetakis, Supervisor
(Department of Computer Science)

Dr. Alexandra Branzan Albu, Departmental Member
(Department of Computer Science)

Dr. Wu-Sheng Lu, Outside Member
(Department of Electrical and Computer Engineering)

Dr. Michael H. Brill, Additional Member
(Datacolor, New Jersey, USA)

ABSTRACT

In computer graphics, the complex phenomenon of color appearance, involving the interaction of light, matter and the human visual system, is modeled by the multiplication of RGB triplets assigned to lights and materials. This efficient heuristic produces plausible images because the triplets assigned to materials usually function as color specifications. *To predict* color, spectral rendering is required, but the $O(n)$ cost of computing reflections with n -dimensional point-sampled spectra is prohibitive for real-time rendering.

Typical spectra are well approximated by m -dimensional linear models, where $m \ll n$, but computing reflections with this representation requires $O(m^2)$ matrix-vector multiplication. A method by Drew and Finlayson [*JOSA A* 20, 7 (2003), 1181-1193], reduces this cost to $O(m)$ by “sharpening” an $n \times m$ orthonormal basis with a linear transformation, so that the new basis vectors are approximately disjoint. If successful, this transformation allows approximated reflections to be computed as the products of coefficients of lights and materials. Finding the $m \times m$ change of basis

matrix requires solving m eigenvector problems, each needing a choice of wavelengths in which to sharpen the corresponding basis vector. These choices, however, are themselves an optimization problem left unaddressed by the method’s authors.

Instead, we pose a single problem, expressing the total approximation error incurred across all wavelengths as the sum of $d \cdot m^2$ squares for some number d , where, depending on the inherent dimensionality of the rendered reflectance spectra, $m \leq d \ll n$, a number that is independent of the number of approximated reflections. This problem may be solved in real time, or nearly, using standard nonlinear optimization algorithms. Results using a variety of reflectance spectra and three standard illuminants yield errors at or close to the best lower bound attained by projection onto the leading m characteristic vectors of the approximated reflections. Measured as CIEDE2000 color differences, a heuristic proxy for image difference, these errors can be made small enough to be likely imperceptible using values of $4 \leq m \leq 9$.

An examination of this problem reveals a hierarchy of simpler, more quickly solved subproblems whose solutions yield, in the typical case, increasingly inaccurate approximations. Analysis of this hierarchy explains why, in general, the lowest approximation error is not attained by simple spectral sharpening, the smallest of these subproblems, unless the spectral power distributions of all light sources in a scene are sufficiently close to constant functions. Using the methods described in this dissertation, spectra can be rendered in real time as the products of m -dimensional vectors of sharp basis coefficients at a cost that is, in a typical application, a negligible fraction above the cost of RGB rendering.

Contents

Supervisory Committee	ii
Abstract	iii
Table of Contents	v
List of Tables	vii
List of Figures	viii
1 Introduction	1
1.1 Major contributions	4
1.2 Overview	5
2 Light sources, reflectance sets and color computations	9
2.1 Light sources	9
2.2 Reflectance sets	10
2.3 Computing color	12
2.3.1 By the standard observer	13
2.3.2 By the RGB heuristic	13
2.3.3 By prefiltered RGB coordinates	14
3 Reducing the dimensionality of spectra	15
3.1 Using standard bases	16
3.1.1 Box functions and point sampling	16
3.1.2 Polynomials	19
3.1.3 Sinusoids	19
3.1.4 Composite models	20
3.1.5 Wavelets	20
3.2 Using bases derived empirically	21

3.2.1	By characteristic vector analysis	21
3.2.2	By other methods	24
4	Spectral sharpening	25
5	Finding the change of basis matrix	28
5.1	For $q \equiv k \equiv 1$, by diagonalizing $\mathbf{B}^T \text{diag}(\mathbf{E})^{-1} \mathbf{B}$	31
5.2	By minimizing the residual $\ \Delta\ _F^2$	51
5.3	By approximate joint diagonalization of $\{\mathbf{R}^{(i)}\}_{i=1}^n$	51
5.4	By sharpening \mathbf{Q}	53
5.5	By minimizing the total error, $\ \Delta\mathbf{c}\ _F^2$	57
6	Results	91
6.1	Accuracy	91
6.1.1	Single light source, direct illumination only ($q \equiv k \equiv 1$)	96
6.1.2	Direct and indirect illumination, $k > 1$	96
6.1.3	Some rendered examples	122
6.2	Real-Time performance	122
7	Conclusion	126
7.1	Summary	126
7.2	Future Work	127
	Appendix	132
	Bibliography	169

List of Tables

Table 1.1	Notation	6
Table 1.2	Symbols	7
Table 5.1	Optimization problems	57
Table 5.2	BFGS $f^{(1)}$ call counts and run times	60
Table 6.1	Relative rendering costs	125

List of Figures

Figure 2.1	Standard illuminants D65, A and F2.	10
Figure 2.2	METACOW image	11
Figure 2.3	METACOW spectra	12
Figure 2.4	CIE 1931 2° Standard Observer color matching functions.	13
Figure 3.1	Variance accounted for by the first five characteristic vectors of Munsell color signals	22
Figure 4.1	First four Munsell/D65 characteristic vectors, before and after sharpening.	25
Figure 5.1	Munsell/D65 $\mathcal{C}^{(1)}$ bases $\mathbf{B}^{(1)}$ and $\mathbf{Q}^{(\mathbf{V})}$	33
Figure 5.2	Munsell/D65 $\mathcal{C}^{(1)}$ bases $\mathbf{B}^{(1)}$ and $\mathbf{Q}^{(5)}$	34
Figure 5.3	Munsell/A $\mathcal{C}^{(1)}$ bases $\mathbf{B}^{(1)}$ and $\mathbf{Q}^{(\mathbf{V})}$	35
Figure 5.4	Munsell/A $\mathcal{C}^{(1)}$ bases $\mathbf{B}^{(1)}$ and $\mathbf{Q}^{(5)}$	36
Figure 5.5	Munsell/F2 $\mathcal{C}^{(1)}$ bases $\mathbf{B}^{(1)}$ and $\mathbf{Q}^{(\mathbf{V})}$	37
Figure 5.6	Munsell/F2 $\mathcal{C}^{(1)}$ bases $\mathbf{B}^{(1)}$ and $\mathbf{Q}^{(5)}$	38
Figure 5.7	Natural Colors/D65 $\mathcal{C}^{(1)}$ bases $\mathbf{B}^{(1)}$ and $\mathbf{Q}^{(\mathbf{V})}$	39
Figure 5.8	Natural Colors/D65 $\mathcal{C}^{(1)}$ bases $\mathbf{B}^{(1)}$ and $\mathbf{Q}^{(5)}$	40
Figure 5.9	Natural Colors/A $\mathcal{C}^{(1)}$ bases $\mathbf{B}^{(1)}$ and $\mathbf{Q}^{(\mathbf{V})}$	41
Figure 5.10	Natural Colors/A $\mathcal{C}^{(1)}$ bases $\mathbf{B}^{(1)}$ and $\mathbf{Q}^{(5)}$	42
Figure 5.11	Natural Colors/F2 $\mathcal{C}^{(1)}$ bases $\mathbf{B}^{(1)}$ and $\mathbf{Q}^{(\mathbf{V})}$	43
Figure 5.12	Natural Colors/F2 $\mathcal{C}^{(1)}$ bases $\mathbf{B}^{(1)}$ and $\mathbf{Q}^{(5)}$	44
Figure 5.13	METACOW/D65 $\mathcal{C}^{(1)}$ bases $\mathbf{B}^{(1)}$ and $\mathbf{Q}^{(\mathbf{V})}$	45
Figure 5.14	METACOW/D65 $\mathcal{C}^{(1)}$ bases $\mathbf{B}^{(1)}$ and $\mathbf{Q}^{(5)}$	46
Figure 5.15	METACOW/A $\mathcal{C}^{(1)}$ bases $\mathbf{B}^{(1)}$ and $\mathbf{Q}^{(\mathbf{V})}$	47
Figure 5.16	METACOW/A $\mathcal{C}^{(1)}$ bases $\mathbf{B}^{(1)}$ and $\mathbf{Q}^{(5)}$	48
Figure 5.17	METACOW/F2 $\mathcal{C}^{(1)}$ bases $\mathbf{B}^{(1)}$ and $\mathbf{Q}^{(\mathbf{V})}$	49
Figure 5.18	METACOW/F2 $\mathcal{C}^{(1)}$ bases $\mathbf{B}^{(1)}$ and $\mathbf{Q}^{(5)}$	50

Figure 5.19	$\mathbf{Q}^{(3)+T} \circ \mathbf{Q}^{(3)}$ and $\mathbf{Q}^{(3)+T}$	53
Figure 5.20	Munsell/(D65, A, F2) $\mathcal{C}^{(1)}$ bases $\mathbf{B}^{(1)}$ and $\mathbf{Q}^{(1)}$	61
Figure 5.21	Munsell/(D65, A, F2) $\mathcal{C}^{(1)}$ bases $\mathbf{B}^{(1)}$ and $\mathbf{Q}^{(5)}$	62
Figure 5.22	Natural Colors/(D65, A, F2) $\mathcal{C}^{(1)}$ bases $\mathbf{B}^{(1)}$ and $\mathbf{Q}^{(1)}$	63
Figure 5.23	Natural Colors/(D65, A, F2) $\mathcal{C}^{(1)}$ bases $\mathbf{B}^{(1)}$ and $\mathbf{Q}^{(5)}$	64
Figure 5.24	METACOW/(D65, A, F2) $\mathcal{C}^{(1)}$ bases $\mathbf{B}^{(1)}$ and $\mathbf{Q}^{(1)}$	65
Figure 5.25	METACOW/(D65, A, F2) $\mathcal{C}^{(1)}$ bases $\mathbf{B}^{(1)}$ and $\mathbf{Q}^{(5)}$	66
Figure 5.26	Munsell/D65 $\mathcal{C}^{(2)}$ bases $\mathbf{B}^{(2)}$ and $\mathbf{Q}^{(1)}$	67
Figure 5.27	Munsell/D65 $\mathcal{C}^{(2)}$ bases $\mathbf{B}^{(2)}$ and $\mathbf{Q}^{(5)}$	68
Figure 5.28	Munsell/A $\mathcal{C}^{(2)}$ bases $\mathbf{B}^{(2)}$ and $\mathbf{Q}^{(1)}$	69
Figure 5.29	Munsell/A $\mathcal{C}^{(2)}$ bases $\mathbf{B}^{(2)}$ and $\mathbf{Q}^{(5)}$	70
Figure 5.30	Munsell/F2 $\mathcal{C}^{(2)}$ bases $\mathbf{B}^{(2)}$ and $\mathbf{Q}^{(1)}$	71
Figure 5.31	Munsell/F2 $\mathcal{C}^{(2)}$ bases $\mathbf{B}^{(2)}$ and $\mathbf{Q}^{(5)}$	72
Figure 5.32	Munsell/(D65, A, F2) $\mathcal{C}^{(2)}$ bases $\mathbf{B}^{(2)}$ and $\mathbf{Q}^{(1)}$	73
Figure 5.33	Munsell/(D65, A, F2) $\mathcal{C}^{(2)}$ bases $\mathbf{B}^{(2)}$ and $\mathbf{Q}^{(5)}$	74
Figure 5.34	Natural Colors/D65 $\mathcal{C}^{(2)}$ bases $\mathbf{B}^{(2)}$ and $\mathbf{Q}^{(1)}$	75
Figure 5.35	Natural Colors/D65 $\mathcal{C}^{(2)}$ bases $\mathbf{B}^{(2)}$ and $\mathbf{Q}^{(5)}$	76
Figure 5.36	Natural Colors/A $\mathcal{C}^{(2)}$ bases $\mathbf{B}^{(2)}$ and $\mathbf{Q}^{(1)}$	77
Figure 5.37	Natural Colors/A $\mathcal{C}^{(2)}$ bases $\mathbf{B}^{(2)}$ and $\mathbf{Q}^{(5)}$	78
Figure 5.38	Natural Colors/F2 $\mathcal{C}^{(2)}$ bases $\mathbf{B}^{(2)}$ and $\mathbf{Q}^{(1)}$	79
Figure 5.39	Natural Colors/F2 $\mathcal{C}^{(2)}$ bases $\mathbf{B}^{(2)}$ and $\mathbf{Q}^{(5)}$	80
Figure 5.40	Natural Colors/(D65, A, F2) $\mathcal{C}^{(2)}$ bases $\mathbf{B}^{(2)}$ and $\mathbf{Q}^{(1)}$	81
Figure 5.41	Natural Colors/(D65, A, F2) $\mathcal{C}^{(2)}$ bases $\mathbf{B}^{(2)}$ and $\mathbf{Q}^{(5)}$	82
Figure 5.42	METACOW/D65 $\mathcal{C}^{(2)}$ bases $\mathbf{B}^{(2)}$ and $\mathbf{Q}^{(1)}$	83
Figure 5.43	METACOW/D65 $\mathcal{C}^{(2)}$ bases $\mathbf{B}^{(2)}$ and $\mathbf{Q}^{(5)}$	84
Figure 5.44	METACOW/A $\mathcal{C}^{(2)}$ bases $\mathbf{B}^{(2)}$ and $\mathbf{Q}^{(1)}$	85
Figure 5.45	METACOW/A $\mathcal{C}^{(2)}$ bases $\mathbf{B}^{(2)}$ and $\mathbf{Q}^{(5)}$	86
Figure 5.46	METACOW/F2 $\mathcal{C}^{(2)}$ bases $\mathbf{B}^{(2)}$ and $\mathbf{Q}^{(1)}$	87
Figure 5.47	METACOW/F2 $\mathcal{C}^{(2)}$ bases $\mathbf{B}^{(2)}$ and $\mathbf{Q}^{(5)}$	88
Figure 5.48	METACOW/(D65, A, F2) $\mathcal{C}^{(2)}$ bases $\mathbf{B}^{(2)}$ and $\mathbf{Q}^{(1)}$	89
Figure 5.49	METACOW/(D65, A, F2) $\mathcal{C}^{(2)}$ bases $\mathbf{B}^{(2)}$ and $\mathbf{Q}^{(5)}$	90
Figure 6.1	Single- and multi-illuminant Munsell $\mathcal{C}^{(1)}$ errors	93
Figure 6.2	Single- and multi-illuminant Natural Colors $\mathcal{C}^{(1)}$ errors	94
Figure 6.3	Single- and multi-illuminant METACOW $\mathcal{C}^{(1)}$ errors	95

Figure 6.4	Worst approximations of $\mathcal{C}^{(1)}$ /D65 color signals by $\mathbf{Q}^{(\mathbf{v})}$. . .	97
Figure 6.5	Worst approximations of $\mathcal{C}^{(1)}$ /D65 color signals by $\mathbf{Q}^{(5)}$. . .	98
Figure 6.6	Worst approximations of $\mathcal{C}^{(1)}$ /illuminant A color signals by $\mathbf{Q}^{(\mathbf{v})}$	99
Figure 6.7	Worst approximations of $\mathcal{C}^{(1)}$ /illuminant A color signals by $\mathbf{Q}^{(5)}$	100
Figure 6.8	Worst approximations of $\mathcal{C}^{(1)}$ /F2 color signals by $\mathbf{Q}^{(\mathbf{v})}$. . .	101
Figure 6.9	Worst approximations of $\mathcal{C}^{(1)}$ /F2 color signals by $\mathbf{Q}^{(5)}$. . .	102
Figure 6.10	Worst approximations of $\mathcal{C}^{(1)}$ /(D65, A, F2) color signals by $\mathbf{Q}^{(1)}$	103
Figure 6.11	Worst approximations of $\mathcal{C}^{(1)}$ /(D65, A, F2) color signals by $\mathbf{Q}^{(5)}$	104
Figure 6.12	Single-illuminant Munsell $\mathcal{C}^{(3)}$ errors incurred by simply sharpening \mathbf{Q}	105
Figure 6.13	Comparison of $\mathbf{Q}^{(5)} \circ \mathbf{Q}^{(5)}$ and $\mathbf{Q}^{(i)} \circ \mathbf{Q}^{(i)}$, $i \in \{1, 4, 6\}$. . .	106
Figure 6.14	Single- and multi-illuminant Munsell $\mathcal{C}^{(3)}$ errors using $\mathbf{Q}^{(i)}$, $i \in \{1, 2, 3, 5\}$	111
Figure 6.15	Single- and multi-illuminant Natural Colors $\mathcal{C}^{(3)}$ errors using $\mathbf{Q}^{(i)}$, $i \in \{1, 2, 3, 5\}$	112
Figure 6.16	Single- and multi-illuminant METACOW $\mathcal{C}^{(3)}$ errors using $\mathbf{Q}^{(i)}$, $i \in \{1, 2, 3, 5\}$	113
Figure 6.17	Worst approximations of $\mathcal{C}^{(3)}$ /D65 color signals by $\mathbf{Q}^{(1)}$. . .	114
Figure 6.18	Worst approximations of $\mathcal{C}^{(3)}$ /D65 color signals by $\mathbf{Q}^{(5)}$. . .	115
Figure 6.19	Worst approximations of $\mathcal{C}^{(3)}$ /illuminant A color signals by $\mathbf{Q}^{(1)}$	116
Figure 6.20	Worst approximations of $\mathcal{C}^{(3)}$ /illuminant A color signals by $\mathbf{Q}^{(5)}$	117
Figure 6.21	Worst approximations of $\mathcal{C}^{(3)}$ /F2 color signals by $\mathbf{Q}^{(1)}$. . .	118
Figure 6.22	Worst approximations of $\mathcal{C}^{(3)}$ /F2 color signals by $\mathbf{Q}^{(5)}$. . .	119
Figure 6.23	Worst approximations of $\mathcal{C}^{(3)}$ /(D65, A, F2) color signals by $\mathbf{Q}^{(1)}$	120
Figure 6.24	Worst approximations of $\mathcal{C}^{(3)}$ /(D65, A, F2) color signals by $\mathbf{Q}^{(5)}$	121
Figure 6.25	Some rendered examples	123

Figure 7.1	Discrete cosines and Munsell $\mathcal{C}^{(1)}$ orthonormal bases \mathbf{B} , $\mathbf{Q}^{(5)}$ and $\mathbf{B}\mathbf{B}^T$	131
Figure A1	Distributions of single-illuminant Munsell $\mathcal{C}^{(1)}$ spectral errors	133
Figure A2	Distributions of single-illuminant Munsell $\mathcal{C}^{(1)}$ color differences	134
Figure A3	Distributions of multi-illuminant Munsell $\mathcal{C}^{(1)}$ spectral errors	135
Figure A4	Distributions of multi-illuminant Munsell $\mathcal{C}^{(1)}$ color differences	136
Figure A5	Distributions of single-illuminant Natural Colors $\mathcal{C}^{(1)}$ spectral errors	137
Figure A6	Distributions of single-illuminant Natural Colors $\mathcal{C}^{(1)}$ color differences	138
Figure A7	Distributions of multi-illuminant Natural Colors $\mathcal{C}^{(1)}$ spectral errors	139
Figure A8	Distributions of multi-illuminant Natural Colors $\mathcal{C}^{(1)}$ color differences	140
Figure A9	Distributions of single-illuminant METACOW $\mathcal{C}^{(1)}$ spectral errors	141
Figure A10	Distributions of single-illuminant METACOW $\mathcal{C}^{(1)}$ color differences	142
Figure A11	Distributions of multi-illuminant METACOW $\mathcal{C}^{(1)}$ spectral errors	143
Figure A12	Distributions of multi-illuminant METACOW $\mathcal{C}^{(1)}$ color differences	144
Figure A13	Distributions of Munsell/D65 $\mathcal{C}^{(3)}$ NRMS spectral errors	145
Figure A14	Distributions of Munsell/D65 $\mathcal{C}^{(3)}$ color differences	146
Figure A15	Distributions of Munsell/A $\mathcal{C}^{(3)}$ NRMS spectral errors	147
Figure A16	Distributions of Munsell/A $\mathcal{C}^{(3)}$ color differences	148
Figure A17	Distributions of Munsell/F2 $\mathcal{C}^{(3)}$ NRMS spectral errors	149
Figure A18	Distributions of Munsell/F2 $\mathcal{C}^{(3)}$ color differences	150
Figure A19	Distributions of Munsell/(D65, A, F2) $\mathcal{C}^{(3)}$ NRMS spectral errors	151
Figure A20	Distributions of Munsell/(D65, A, F2) $\mathcal{C}^{(3)}$ color differences	152
Figure A21	Distributions of Natural Colors/D65 $\mathcal{C}^{(3)}$ NRMS spectral errors	153
Figure A22	Distributions of Natural Colors/D65 $\mathcal{C}^{(3)}$ color differences	154
Figure A23	Distributions of Natural Colors/A $\mathcal{C}^{(3)}$ NRMS spectral errors	155

Figure A24	Distributions of Natural Colors/A $\mathcal{C}^{(3)}$ color differences . . .	156
Figure A25	Distributions of Natural Colors/F2 $\mathcal{C}^{(3)}$ NRMS spectral errors	157
Figure A26	Distributions of Natural Colors/F2 $\mathcal{C}^{(3)}$ color differences . . .	158
Figure A27	Distributions of Natural Colors/(D65, A, F2) $\mathcal{C}^{(3)}$ NRMS spectral errors	159
Figure A28	Distributions of Natural Colors/(D65, A, F2) $\mathcal{C}^{(3)}$ color differences	160
Figure A29	Distributions of METACOW/D65 $\mathcal{C}^{(3)}$ NRMS spectral errors	161
Figure A30	Distributions of METACOW/D65 $\mathcal{C}^{(3)}$ color differences . . .	162
Figure A31	Distributions of METACOW/A $\mathcal{C}^{(3)}$ NRMS spectral errors	163
Figure A32	Distributions of METACOW/A $\mathcal{C}^{(3)}$ color differences . . .	164
Figure A33	Distributions of METACOW/F2 $\mathcal{C}^{(3)}$ NRMS spectral errors	165
Figure A34	Distributions of METACOW/F2 $\mathcal{C}^{(3)}$ color differences . . .	166
Figure A35	Distributions of METACOW/(D65, A, F2) $\mathcal{C}^{(3)}$ NRMS spectral errors	167
Figure A36	Distributions of METACOW/(D65, A, F2) $\mathcal{C}^{(3)}$ color differences	168

Chapter 1

Introduction

Color appearance is an observer's perception of the interaction of light with an object's material. A prediction or an explanation of color requires models of the incident light, the material's optical properties and the observer's visual system. First-order models are constructed with a light source's spectral power distribution (SPD), which is a measure of the radiative power present at each wavelength within the visual range, and a material's reflectance spectrum, which describes the fraction of incident light that the material reflects as a function of wavelength¹. The simplest model of a visual system consists of some number of spectral sensitivity functions, one for each type of its photodetectors. A color is formed by the inner products of these and an observed SPD, called a *color signal*, which is the pointwise product of a reflectance spectrum and a light source's SPD. In the trichromatic human visual system, at normal, or *photopic*, light levels, a color has three dimensions, one for each type of cone cell on the retina. In linear algebraic terms, a color in this simple model is the projection of an infinite-dimensional color signal onto the three-dimensional subspace spanned by the cone cells' spectral sensitivities.

Common practice in computer graphics omits an explicit model of the observer and represents lights and materials as RGB triplets, which are multiplied together to compute a color. This model of color formation is often described as approximating spectra with coarse samplings at red, green and blue wavelengths. Functionally, however, the triplet assigned to a material is really less an optical property than a *color specification*, the brightness of which is modulated by a light source, typically a

¹Although the models and methods described here can apply equally well to light that is transmitted by a material, for simplicity the discussion focuses on reflection. For convenience, the term *spectrum* is used for SPDs and reflectance spectra.

scaled white. This simple and efficient model suffices because color prediction is not the usual goal in computer graphics, which is the creation of images that are plausible, pleasing or consistent with a specification. Colors in these images can be accurate as well, if the spectra of the corresponding light sources or materials are sufficiently close to constant functions [13]. In global illumination algorithms, however, because the SPDs of reflections by chromatic materials are usually far from constant functions, colors resulting from indirect illumination can be inaccurate enough to be implausible.

There are, however, phenomena affecting color appearance that RGB triplets cannot model at any level of accuracy. The information lost when color signals are projected onto a three-dimensional subspace can cause different SPDs, known as *metamers*, to be perceived as matching colors. Two objects can thus have colors that match when illuminated by one light source but differ when illuminated by another. This mundane color appearance phenomenon cannot happen in the RGB model. Two triplets that are equal when multiplied by a third with non-zero components will remain equal when multiplied by any other. Nor can RGB triplets model iridescence, an effect of wavelength-dependent modulation of reflected light by coherent interference [10, 96, 135, 62].

Its inadequacy in color appearance prediction is not the RGB model's only limitation. When the subject of interest is spectral information, it is not even relevant. This is plainly the case when wavelengths outside the visual range are involved, as in flight simulators that integrate radar and infrared sensors with visual displays [141]. Other examples include the rendering of 3D hyperspectral models of art or heritage objects not only for appearance prediction [42, 19], but also for the analysis of material properties [25, 95], or of biological specimens for the simulation of color appearance in non-human visual systems [94]. These applications extend spectral information to ultraviolet and infrared wavelengths, as do some simulations of the optical properties of biological tissues for study and the diagnosis of disease [29]. In simulations of air- or spaceborne remote imaging, models of the interaction of light with forests or crops can determine the biochemical state of the plants they contain by inverting hyperspectral images [47]. Spectral rendering is also used in the design of lighting [65] and optical instruments [131], as well as in the study of human vision [45] and computer vision [44]. In scientific visualization applications the higher dimensionality of spectra can convey information with false color images by exploiting, for example, metamerism in volume rendering [9] or iridescence in the visualization of photoelasticity by a virtual polariscope [18].

When spectra are represented as n -dimensional vectors of point samples, reflections are computed by componentwise vector multiplication with complexity $O(n)$ and mapped to colors with a $3 \times n$ linear transformation. Because typical spectra are highly correlated, they can be approximated with sufficient accuracy by a linear model in which they are represented as m -dimensional vectors of coefficients of an orthonormal basis obtained by a characteristic vector analysis. This approach, however, is practical only for small values of $m \ll n$, as computing reflections involves matrix-vector multiplication with complexity $O(m^2)$.

A method proposed by Drew and Finlayson reduces this cost by using a change of coordinates “to sharpen” an orthonormal basis so that its vectors are approximately disjoint, or complementary [49]. If the transformation is successful, approximated reflections may be computed with sufficient accuracy and complexity $O(m)$ by componentwise multiplication of transformed basis coefficients of SPDs and reflectance spectra, from which colors are obtained with a $3 \times m$ linear transformation. Drew and Finlayson find the change of basis matrix by solving m eigenvector problems, each of which requires the user to specify an interval of wavelengths within which the corresponding basis vector is sharpened. Finding the wavelengths that result in the sharpest basis vectors, however, is itself a problem left unaddressed by the authors, nor do they pose or answer the question of whether maximizing sharpness does in fact minimize approximation error.

In this work we show that the problem of choosing wavelength intervals is avoided by simply minimizing the total error incurred across all wavelengths, expressed as the sum of $d \cdot m^2$ squares (where $m \leq d \ll n$), a number that is independent of the number of approximated spectra and small enough to be solved in milliseconds or less. An examination of this expression shows that while sharpening of the basis does play a role in minimizing the error, it does not by itself yield the optimal transformation, unless all light source SPDs are sufficiently close to constant functions. Instead, a nonlinear optimization can minimize the error explicitly, to a level close to the theoretical minimum attained by projection onto the approximated spectra’s leading m characteristic vectors. For our test data consisting of three standard illuminants and a variety of reflectance spectra, all represented as 61-dimensional vectors, bases of only four to nine dimensions yield sufficient accuracy, even when second and third reflections are included. To understand the role of spectral sharpening in this optimization, we also discuss and present results obtained from a hierarchy of subproblems yielding approximations that, while less accurate, may be accurate enough for some applications

needing fast solutions when unknown lighting conditions change in real time. The subproblem of spectral sharpening, the smallest though least accurate of these, can be solved in microseconds.

1.1 Major contributions

Summarized, the major contributions of this work are:

- A storage-efficient representation of spectra that allows real-time spectral rendering at negligible cost above that of RGB rendering, with an accuracy depending only on m , the representation's dimensionality.
- A mathematical exposition and analysis of this representation and its relationship to spectral sharpening, the smallest and least optimal in a hierarchy of subproblems derived from an expression of the approximation error to be minimized.
- An expression of the approximation error as the sum of $d \cdot m^2$ squares, where $d \ll n$ is the reflectance spectra's inherent dimensionality and $4 \leq m \leq 9$. The complexity of this expression is independent of the number of approximated SPDs and small enough to be optimized in real time or nearly.
- Suggestions for methods of solving this optimization problem and those in the hierarchy of smaller subproblems, taken from existing methods in nonlinear optimization, approximate joint diagonalization and factor analysis.
- A detailed reporting of the efficiency of these optimizations and the accuracy attained by their solutions for first, second and third reflections using three standard illuminants and three diverse sets of reflectance spectra. Imperceptible approximation errors, attained with 4- to 9-dimensional representations, are in all cases at or very near the theoretical lower bound attained by projection onto the best-fitting orthonormal basis of the same dimensionality obtained by a characteristic vector analysis.

1.2 Overview

This dissertation is organized into seven chapters, including this one, and some supplementary material. Notation used throughout is explained in Table 1.1 and symbols are defined in Table 1.2. In the next chapter the reflectance spectra and illuminant SPDs used to test the methods presented in later chapters are described. A simple model of color formation from SPDs is also described and contrasted with the usual practice in computer graphics of multiplying RGB triplets. In Chapter 3 various ways to reduce the dimensionality of spectra are reviewed. These techniques, described in the computer graphics, color science or hyperspectral imaging literature, seek to reduce the cost of computation or storage, or possibly both, as does Drew's and Finlayson's method of spectral sharpening, which is presented in Chapter 4. At the heart of this method is an optimization problem, the solution to which is a change of basis matrix, denoted \mathbf{T} , representing a linear transformation to a "sharp" coordinate system in which reflections may be computed by multiplying coefficients of a transformed lower-dimensional orthonormal basis. Different versions of this optimization problem and methods for solving them are the subject of Chapter 5. Derived from an unconstrained, nonlinear, least squares problem whose solution minimizes a residual error, a hierarchy of smaller subproblems is described. These are useful both for finding approximate solutions more quickly and for understanding the role of spectral sharpening, and its insufficiency, in minimizing the residual error. Real-time rendering performance and the accuracy attained by these methods, measured as spectral errors and resulting color differences, are reported in Chapter 6 and the supplementary material in the appendix. In closing, some conclusions are drawn and possible future directions are discussed in Chapter 7.

Table 1.1: Notation

n, λ	scalar (not boldface)
\mathcal{X}	set (calligraphic, not boldface)
\mathbf{X}, \mathbf{x}	matrix or column vector (boldface)
\mathcal{X}	matrix with columns populated by the elements of \mathcal{X} (boldface)
\mathbf{X}^T ,	transpose of \mathbf{X}
$\mathbf{X}^{-1}, \mathbf{X}^{-T}$	inverse of \mathbf{X} and its transpose
$\mathbf{X}^+, \mathbf{X}^{+T}$	Moore-Penrose pseudoinverse of \mathbf{X} and its transpose
\mathbf{X}_{ij}	row i , column j element of \mathbf{X}
\mathbf{X}_i	column (element) i of the matrix (vector) \mathbf{X}
$\mathbf{X}_{i:j}$	columns i through j of \mathbf{X}
\mathbf{X}^i	\mathbf{X} raised to the i -th power
$\ \mathbf{X}\ _F$	the Frobenius norm of \mathbf{X} : $\sqrt{\sum_{i,j} \mathbf{X}_{ij}^2}$
$\sqrt{\mathbf{X}}$	matrix square root, $\sqrt{\mathbf{X}}\sqrt{\mathbf{X}} = \mathbf{X}$, for positive definite \mathbf{X}
$ \mathbf{X} $	matrix containing the absolute values of the elements of \mathbf{X}
$\mathbf{X}^{(i)}, \mathbf{X}^{(i)T}$	a particular matrix \mathbf{X} labeled i and its transpose
$\langle \mathbf{x}, \mathbf{y} \rangle$	inner product of vectors \mathbf{x} and \mathbf{y}
\mathbf{I}	the identity matrix
$\mathbf{1}$	a vector of ones
$\mathbf{0}$	the zero vector
$\text{diag}(\mathbf{x})$	matrix formed with vector \mathbf{x} on the diagonal and zeros elsewhere
$\text{diag}(\mathbf{X})$	vector formed by the diagonal entries of the square matrix \mathbf{X}
$\text{ddiag}(\mathbf{X})$	$\text{diag}(\text{diag}(\mathbf{X}))$, for square matrix \mathbf{X}
$\text{off}(\mathbf{X})$	$\mathbf{X} - \text{ddiag}(\mathbf{X})$, for square matrix \mathbf{X}
$\text{tr}(\mathbf{X})$	the trace of the matrix \mathbf{X}
$\mathbf{X} \circ \mathbf{Y}$	the Hadamard (componentwise) product of \mathbf{X} and \mathbf{Y}
$\mathbf{X} \otimes \mathbf{Y}$	the Kronecker product of \mathbf{X} and \mathbf{Y}
$\mathbf{X} * \mathbf{Y}$	the Khatri-Rao (or columnwise Kronecker) product of \mathbf{X} and \mathbf{Y}

Table 1.2: Symbols

scalars	
λ	wavelength (in nanometers, nm)
n	dimensionality of point-sampled spectra
m	dimensionality (i.e., rank) of the basis \mathbf{B}
p, q, s	sizes of sets \mathcal{S} , \mathcal{E} and $\mathcal{C}^{(k)}$
r, g, b	red, green and blue color components
k	number of reflections used to construct $\mathcal{C}^{(k)}$
sets	
\mathcal{S}	point-sampled reflectance spectra
\mathcal{E}	point-sampled light source SPDs
$\mathcal{C}^{(k)}$	point-sampled SPDs formed by up to k reflections
\mathcal{L}	set of subsets of wavelengths $\{400, 405, \dots, 700\}$ (in nm)
vectors and matrices	
\mathbf{S}	$n \times 1$ point-sampled reflectance spectrum $S(\lambda)$
\mathbf{E}	$n \times 1$ point-sampled light source SPD $E(\lambda)$
\mathbf{C}	$n \times 1$ point-sampled reflected SPD $C(\lambda) \equiv E(\lambda)S(\lambda)$
\mathbf{M}	$3 \times n$ SPD to RGB transformation matrix
\mathbf{b}	3×1 linear color
$\mathbf{U}^{(\mathbf{X})}\mathbf{\Sigma}^{(\mathbf{X})}\mathbf{V}^{(\mathbf{X})\text{T}}$	singular value decomposition (SVD) of the matrix \mathbf{X}
$\boldsymbol{\sigma}^{(\mathbf{X})}$	singular values of the matrix \mathbf{X} , $\text{diag}(\mathbf{\Sigma}^{(\mathbf{X})})$
$\hat{\mathbf{U}}^{(\mathbf{X})} \equiv \mathbf{U}^{(\mathbf{X})}\mathbf{\Sigma}^{(\mathbf{X})}$	weighted left singular vectors of the matrix \mathbf{X}
\mathcal{S}, \mathcal{E}	$n \times p$ and $n \times q$ matrices of elements of \mathcal{S} and \mathcal{E}
$\mathcal{C}^{(k)}$	$n \times s$ matrix of elements of $\mathcal{C}^{(k)}$
\mathcal{C}	$\mathcal{C}^{(1)}$, unless otherwise specified
$\mathbf{B}^{(k)}$	$n \times m$ orthonormal basis for $\mathcal{C}^{(k)}$
\mathbf{B}	$\mathbf{B}^{(2)}$, unless otherwise specified
$\mathbf{R}^{(i)}$	$\mathbf{B}^{\text{T}} \text{diag}(\mathbf{S}^{(i)})\mathbf{B}$, $m \times m$ operator for reflection by $\mathbf{S}^{(i)}$
$\mathbf{s}, \mathbf{e}, \mathbf{c}$	$m \times 1$ basis coefficients $\mathbf{B}^{\text{T}}\mathbf{S}$, $\mathbf{B}^{\text{T}}\mathbf{E}$, $\mathbf{B}^{\text{T}}\mathbf{C}$
$\boldsymbol{\mathcal{S}}, \boldsymbol{\mathcal{E}}$	$m \times p$ and $m \times q$ matrices of basis coefficients $\mathbf{B}^{\text{T}}\mathcal{S}$ and $\mathbf{B}^{\text{T}}\mathcal{E}$
$\mathbf{c}^{(k)}$	$m \times s$ matrix of basis coefficients $\mathbf{B}^{\text{T}}\mathcal{C}^{(k)}$
\mathbf{c}	$\mathbf{c}^{(1)}$, unless otherwise specified

Table 1.2: (continued)

vectors and matrices (continued)	
T	$m \times m$ change of basis sharpening matrix
Q	$n \times m$ matrix of sharp basis vectors \mathbf{BT}
$\tilde{\mathbf{s}}, \tilde{\mathbf{e}}, \tilde{\mathbf{c}}$	$m \times 1$ sharp basis coefficients of \mathbf{S} , \mathbf{E} and \mathbf{C}
$\tilde{\mathbf{d}}, \tilde{\mathbf{e}}$	$m \times p$ and $m \times q$ matrices of basis coefficients $\mathbf{Q}^+\mathcal{S}$ and $\mathbf{Q}^+\mathcal{E}$
$\tilde{\mathbf{M}}$	$3 \times m$ sharp basis coefficients to RGB transformation matrix
$\tilde{\mathbf{u}}^{(\mathbf{x})} \equiv \mathbf{Q}^+\hat{\mathbf{U}}^{(\mathbf{x})}$	sharp basis coefficients of $\hat{\mathbf{U}}^{(\mathbf{x})}$
Δ	$n \cdot m \times m$ matrix of n vertical blocks of residuals: $\Delta^{(i)} \equiv \mathbf{B}^T \text{diag}(\hat{\mathbf{U}}_i^{(\mathcal{S})})\mathbf{B} - \mathbf{T} \text{diag}(\tilde{\mathbf{u}}_i^{(\mathcal{S})})\mathbf{T}^{-1}, i = 1, \dots, n$

Chapter 2

Light sources, reflectance sets and color computations

2.1 Light sources

In classical optics, *visible light*, or simply *light*, is electromagnetic radiation consisting of superposed waves, propagating in space through the electromagnetic field, with frequencies corresponding to wavelengths within the visible range, approximately 400 to 700 nanometers (nm). In quantum optics, these waves are viewed as discrete photons with energies and in quantities respectively proportional to the waves' frequencies and intensities. The energy content of light radiated by a source, as a function of wavelength λ , can be characterized in both models by a spectral power distribution of the radiant exitance, a ratio measured in units of watts per square meter per nanometer ($\text{Wm}^{-2}\text{nm}^{-1}$):

$$M(\lambda) \equiv \frac{\Phi}{A\Delta\lambda},$$

where Φ is radiant flux, A is the area through which the flux passes and $\Delta\lambda$ is small wavelength interval [67, Chapter 13]. To give a unitless relative spectral power distribution, referred to hereafter simply as an SPD, $M(\lambda)$ is normalized by its value at a particular wavelength λ_0 :

$$E(\lambda) \equiv \frac{M(\lambda)}{M(\lambda_0)}.$$

For use in colorimetry, the International Commission on Illumination (or CIE for

Commission internationale de l'éclairage, its official French name) publishes synthetic SPDs known as the *CIE standard illuminants* [177, 34, 81], each of which is intended to resemble the SPD of a real source of light, with $\lambda_0 = 560$ nm, the approximate wavelength at which human photopic vision is most sensitive. The methods described in this work are tested with three of these, chosen as representatives of commonplace sources of light: D65, a daylight simulator; illuminant A, representing the incandescent light emitted by a typical, domestic, tungsten-filament lamp; and F2, from the CIE's series of fluorescent lamp SPDs. These three are chosen also for their familiarity to the color science community and for their dissimilar features as mathematical functions, shown plotted in Fig. 2.1. D65 is relatively smooth and close to constant; illuminant A is smooth, but far from constant; and F2 has four peaks superimposed on a smooth, varying background, three of which are very sharp.

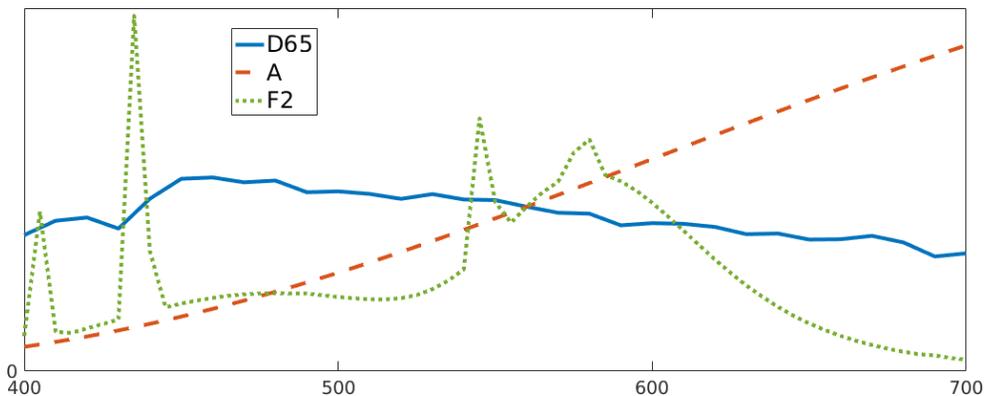


Figure 2.1: Standard illuminants D65, A and F2.

2.2 Reflectance sets

A reflectance spectrum $S(\lambda)$, ranging in value between zero and one¹, gives the fraction of incident light that a material reflects as a function of wavelength, a consequence of its molecular structure [173, 162, 130]. The energy states of molecules are changed in discrete units by absorbing incident photons of corresponding energies, resulting in electronic transitions and lower energy molecular vibrations and rotations. The number of these states, and the exponentially larger number of possible states of lower energy resulting from their interactions in bulk matter, broaden the peaks in absorption spectra corresponding to electronic transitions [115, 75]. This broadening

¹This does not hold for fluorescent or iridescent materials, which are not considered here.

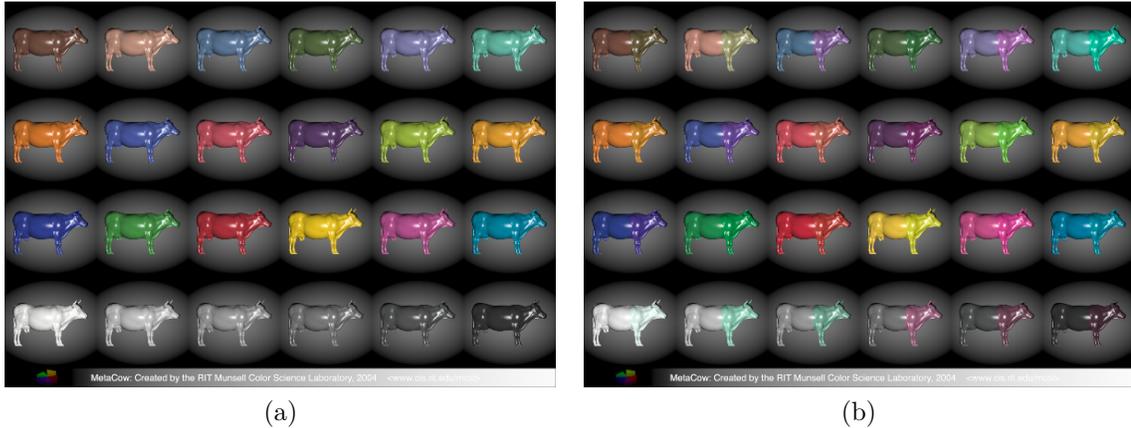


Figure 2.2: METACOW image: a) illuminated by D65; b) illuminated by Illuminant A and chromatically adapted to D65.

accounts for the smoothness of reflectance spectra, which have been characterized as low-bandwidth functions. A set of color signals formed with 348 daylight SPDs and 1,695 reflectance spectra of plant materials, for example, was found by Bonnardel and Maloney to have a frequency limit of approximately 6 cycles/300 nm and energy content of 97% or more in frequencies not exceeding 2 cycles/300 nm [12].

The reflectance spectra used in this work as test cases, in combination with illuminants D65, A and F2, come from three sets, chosen for their different characteristics. The largest of these contains 1,269 relatively smooth spectra of painted chips from the Munsell Matte collection [128], as measured by the University of Eastern Finland Spectral Color Research group² [98]. These spectra, familiar to color scientists, are designed to be seen as colors occupying a large volume in the gamut of perceivable colors, ordered in an approximately uniform distribution according to hue, value and chroma. The same group’s Natural Colors, a set of 218 spectra of plant materials, is chosen for its many saturated colors resulting from spectra that are significantly less smooth than those of the Munsell painted chips. The third set contains the 48 spectra of the METACOW image³, shown in Fig. 2.2, which was offered to the color science community by Fairchild and Johnson as a spectral test target [54]. Half of these are measured from the well known ColorChecker[®] calibration target⁴ [121]. The other half are metamers constructed to be perceived as colors that match when illuminated by D65 but differ significantly when reflecting illuminant A. Most of these unrealistic

²<http://www.uef.fi/web/spectral/-spectral-database>

³http://www.rit.edu/cos/colorscience/rc_db_metacow.php

⁴<http://www.xrite.com/categories/calibration-profiling/colorchecker-classic>

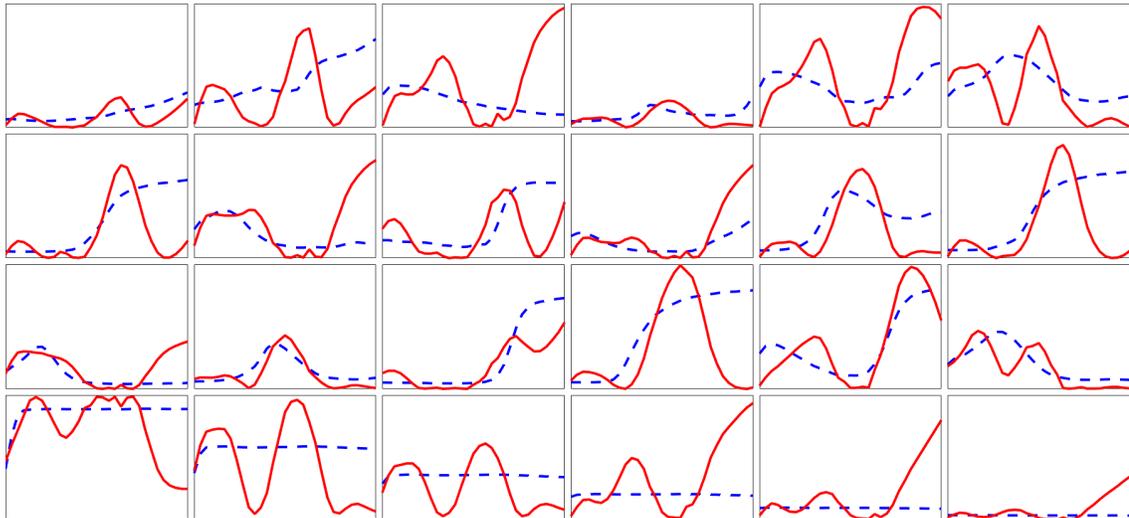


Figure 2.3: METACOW spectra. Measured ColorChecker[®] spectra (rear portions of cows in the METACOW image) are plotted with dashed (blue) lines. Synthetic metamer spectra (front portions) are shown with solid (red) lines. Wavelength and reflectance axes respectively range from 400 to 700 nm and from 0 to 1.

synthetic spectra, shown in Fig. 2.3, are sharper than the sharpest Natural Colors spectra.

2.3 Computing color

A color signal $C(\lambda) \equiv S(\lambda)E(\lambda)$ is the pointwise product of a reflectance spectrum and an illuminant's SPD. A sequence of reflections by k surfaces yields a color signal computed in the obvious way:

$$C(\lambda) \equiv S_k(\lambda) \cdots S_1(\lambda)E(\lambda).$$

When represented as n -dimensional vectors of point samples, spectra can be computed, manipulated and understood using theory and methods from linear algebra and matrix analysis [163, 118]. Sampled every 5 nm uniformly between 400 and 700 nm, a resolution high enough to capture the narrow peaks of fluorescent light source SPDs [104], our spectra are 61-dimensional vectors, and a color signal $\mathbf{C} \in \mathbb{R}^{61}$ is thus

$$\mathbf{C} \equiv \text{diag}(\mathbf{S})\mathbf{E} \equiv \mathbf{S} \circ \mathbf{E}.$$

2.3.1 By the standard observer

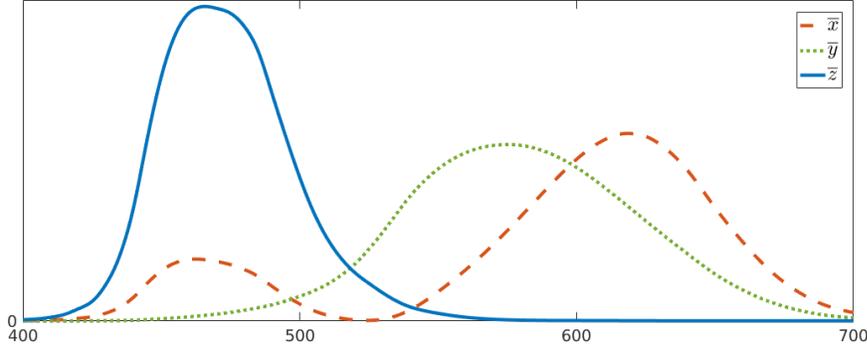


Figure 2.4: CIE 1931 2° Standard Observer color matching functions.

In a simplified model of color formation, a color signal \mathbf{C} is mapped to a color $\mathbf{b}^{(\mathbf{C})} \equiv (r, g, b)^{(\mathbf{C})\text{T}}$ by the linear transformation expressed by a $3 \times n$ matrix \mathbf{M} , the product of three matrices: 1) a $3 \times n$ matrix $\mathbf{M}^{(SPD \rightarrow XYZ)}$ of CIE 1931 2° Standard Observer color matching functions (see Fig. 2.4) mapping SPDs to XYZ coordinates [34, 5]; 2) a 3×3 matrix $\mathbf{M}^{(ca)}$ expressing a chromatic adaptation [53]; and 3) a 3×3 matrix $\mathbf{M}^{(XYZ \rightarrow RGB)}$, which maps adapted XYZ coordinates to an RGB color space, such as Rec. 709, on which the sRGB standard is based [3]:

$$\mathbf{b}^{(\mathbf{C})} \equiv (r, g, b)^{(\mathbf{C})\text{T}} = \mathbf{M}^{(XYZ \rightarrow RGB)} \mathbf{M}^{(ca)} \mathbf{M}^{(SPD \rightarrow XYZ)} (\mathbf{S} \circ \mathbf{E}) \equiv \mathbf{M} \mathbf{C}.$$

In the typical case, the linear color $\mathbf{b}^{(\mathbf{C})}$ is then nonlinearly transformed for display by tone mapping [148], gamut mapping [125] and gamma correction [3].

2.3.2 By the RGB heuristic⁵

It is common practice in computer graphics to compute colors by multiplying RGB triplets assigned to materials and lights. This is equivalent to

$$\mathbf{b}^{(\mathbf{C})'} \equiv (r, g, b)^{(\mathbf{S})\text{T}} \circ (r, g, b)^{(\mathbf{E})\text{T}} = \mathbf{M}(\mathbf{S} \circ \mathbf{1}) \circ \mathbf{M}\mathbf{E} = \mathbf{M}\mathbf{S} \circ \mathbf{M}\mathbf{E},$$

where $\mathbf{1}$ is a constant light source SPD of equal energy, otherwise known as standard illuminant E. As observed in [116], “there is no *mathematical* reason” to expect $\mathbf{b}^{(\mathbf{C})}'$ to equal $\mathbf{b}^{(\mathbf{C})}$, “even approximately.”

⁵This term is borrowed from [116].

In computer graphics, however, usually $(r, g, b)^{(\mathbf{S})\text{T}} \equiv \frac{1}{\alpha}(r, g, b)^{(\mathbf{C})\text{T}}$ and $(r, g, b)^{(\mathbf{E})\text{T}} \equiv \alpha\mathbf{1}$, for some constant α . In this case, $\mathbf{b}^{(\mathbf{C})'}$ and $\mathbf{b}^{(\mathbf{C})}$ are equal, and $(r, g, b)^{(\mathbf{S})\text{T}}$ is a color *specification*, not a coarse spectral sampling of \mathbf{S} , as is often claimed. A light source with equal red, green and blue components simply amplifies or attenuates the brightness of the color specified by the RGB triplet assigned to the material. However, when in fact $\mathbf{E} \equiv \alpha\mathbf{1}$ and $\mathbf{ME} \equiv \alpha\mathbf{1}$, because the SPD $\mathbf{1}$ maps to $\mathbf{1}$ in the color space chosen for display⁶, then these two interpretations of RGB triplets coincide. This observation is implied by Borges’s analysis showing that the error incurred by the RGB heuristic is bounded by the product of the variation in $S(\lambda)$ and $E(\lambda)$ [13].

2.3.3 By prefiltered RGB coordinates

Setting the color $\mathbf{b}^{(\mathbf{E})}$ of the illuminant \mathbf{E} equal to $\mathbf{1}$ and explicitly assigning color specifications to materials as reflectance properties is the essence of the “prefiltering” method described by Ward and Eydelberg-Vileshin, which seeks to make RGB rendering more accurate [172]. When a scene has a “dominant illuminant” it is assigned the color $\mathbf{1}$, and materials are assigned the corresponding “prefiltered” colors $\mathbf{b}^{(\mathbf{C})}$. Colors of materials illuminated exclusively by this illuminant are thus computed correctly. Colors resulting from any indirect illumination, however, are not, nor are colors resulting from illumination by any other light sources in the scene, the colors of which are scaled inversely by the dominant illuminant’s original colors. Of course, if all illuminants in a scene make equal contributions, then all computed colors will be inaccurate, no matter which illuminant is arbitrarily designated as dominant.

⁶In the typical case, when rendering with Rec. 709 triplets, this condition will not be satisfied, since D65 (the color space’s white point illuminant) maps to $\mathbf{1}$, not illuminant \mathbf{E} .

Chapter 3

Reducing the dimensionality of spectra

Rendering with spectra represented as n -dimensional vectors is straightforward and similar to RGB rendering, requiring n multiplications at each reflection, instead of three, and a $3 \times n$ matrix-vector multiplication at each pixel to obtain a color from a color signal. However, if $n \equiv 61$, to accommodate narrow spectral peaks when rendering with fluorescent lights, this cost may be prohibitive. It would preclude real-time rendering for all but the simplest scenes. To reduce this cost, various alternative representations have been proposed in the computer graphics literature. Researchers in hyperspectral imaging have also proposed ways to reduce the dimensionality of spectra, motivated by a desire to reduce the storage and transmission costs of data generated in remote sensing and medical imaging applications. These methods are potentially useful in computer graphics as well, since the rate-limiting factor in rendering applications is most often memory bandwidth, not processing power. That is, additional computation required to multiply or expand spectra represented compactly with only a few parameters can often be easily offset by efficiencies gained in memory access. Our goal, however, is a representation that is efficient for both storage and computation.

Closed-form expressions depending on a small number of parameters exist for many typical light source SPDs. Some of these, like the emission lines of phosphors or the distribution of blackbody radiation described by Planck's law, which depends only on temperature, can be derived from models of the underlying physical processes [109]. Others, such as the CIE's illuminant A and its D series, including D65,

have been defined by continuous functions of wavelength [34]. However, reflectance spectra, and therefore color signals, are not generally available in closed form, as the physical processes involved are not easily modeled. Instead, point-sampled spectra have been represented in fewer dimensions as vectors of coefficients used to form linear combinations of $m \ll n$ basis vectors¹, either chosen from a standard mathematical family, such as low-order polynomials or sinusoids, or derived empirically by methods such as characteristic vector analysis. In general, a spectrum \mathbf{X} is approximated by its projection onto a basis \mathbf{B} as $\mathbf{X} \approx \mathbf{B}\mathbf{B}^+\mathbf{X} \equiv \mathbf{B}\mathbf{x}$, where \mathbf{x} is a vector of basis coefficients, and $\mathbf{B}^+ \equiv (\mathbf{B}^T\mathbf{B})^{-1}\mathbf{B}^T$ is the Moore-Penrose pseudoinverse of \mathbf{B} . Of course, if \mathbf{B} is orthonormal (i.e., $\mathbf{B}^T\mathbf{B} = \mathbf{I}$), then $\mathbf{B}^+ = \mathbf{B}^T$.

3.1 Using standard bases

3.1.1 Box functions and point sampling

Representing a spectrum as a linear combination of box functions is similar to point sampling, but with coefficients computed by weighted averages of spectral values within corresponding wavelength ranges instead of by values at specific wavelengths. In the limit, the two representations coincide as the ranges' widths approach zero. That is, a vector of point samples may be interpreted as coefficients of a basis of boxes of infinitesimal width. For this reason, they are considered here as equivalent, since the choices they involve, with the exception of box width, are similar, namely, their number and placement in the spectrum. Three types of strategies for choosing wavelengths have been described: predetermined and static, dynamically modified at each reflection or chosen stochastically in Monte Carlo methods of simulating light transport.

Meyer describes a procedure for using Gaussian quadrature to select static sampled wavelengths that facilitate numerical integration with a set of color opponent fundamental curves formed by a linear transformation of the 1931 CIE standard observer curves [123]. Testing his results by ColorChecker[®]/illuminant C color differences, he determines that an efficient basis consists of four wavelengths corresponding approximately to the peaks in his color opponent curves. An advantage of this representation is efficient, analytical computation of tristimulus coordinates from spectra,

¹Nonlinear methods of reducing the dimensionality of spectra have also been described, but these are outside the scope of this work. See, for example, [63, 22, 39, 57, 103, 6].

but arbitrary spectra are not well represented with such a low sampling resolution. The author mentions the necessity of extending his technique in the future to accommodate “spikey” spectra, but this work has not yet appeared.

Hall finds the wavelength ranges of a static basis of non-overlapping box functions by minimizing the resulting color differences of the ColorChecker[®] target reflecting four standard illuminants [71, 72]. He reports optimal wavelength choices for dimensionalities of $1 \leq m \leq 10$, yielding root mean square (RMS) $L^*u^*v^*$ color differences of approximately 50 to less than 0.5 for first through fourth reflections of illuminant C and three D illuminants by the spectra of the ColorChecker[®] and four metals.

Zeghers et al. describe a method of selecting $m \ll n$ sampled wavelengths that is based on spectral error, not color difference [181]. Static and dynamic versions of their algorithm select sampled wavelengths indirectly by not sampling those whose omission yields an approximation error within a user-specified bound. Approximated spectra are constructed by linear interpolation after removing candidate wavelengths and adjusting neighboring spectral values to preserve the area under the spectral curve. In the static version these wavelengths are selected in a pre-process by decimating the SPDs of all light sources in a scene. In the dynamic version the selection process is repeated at every reflection by applying it to the reflected SPD. The authors report results obtained with both methods by measuring mean color difference of the pixels in an image rendered with four different light sources, including a fluorescent. The authors claim that the color differences resulting from both methods are not noticeable when as few as six wavelengths are selected. The usefulness of the dynamic version, however, is questionable, as the cost of resampling at every reflection is $O(m \cdot n)$.

Rougeron and Péroche represent every spectrum with its own non-overlapping and contiguous box functions, the number, positions and widths of which are chosen adaptively at every reflection to yield an approximation error within bounds specified in advance by a user [152]. When this error, measured as the norm of an XYZ tristimulus value, exceeds the upper bound, one or more boxes are subdivided until the new approximation’s error reaches a tolerated level. Errors below the lower bound trigger a merging of boxes. Testing their method with the 24 ColorChecker[®] spectra and standard illuminants C, D65 and F2, they find that reasonable approximations of first, second and third reflections respectively require, on average, approximately 8 to 10 (depending on the illuminant), 7 and 5 box functions, and that approximately half of all reflections result in splitting or merging of boxes. Compared to a regular 5-nm point sampling, the overhead required by their method results in slightly longer run

times, but these are more than offset by faster computation of tristimulus coordinates at each pixel resulting from coarser samplings. In a more modern context, however, this may not be the case, as the divergent code paths required to adjust boxes differently on reflections would likely reduce the efficiency gains of parallel computation, especially when rendering on a GPU. A limitation of their method is that it can accommodate only a single light source in a scene. Another is that since the XYZ color space is not perceptually uniform, its use to measure error is questionable, albeit practical, as the cost, while rendering, of the nonlinear computation of error in a perceptually uniform space would be prohibitive. Iehl and Péroche, using the same adaptive representation, address this shortcoming in an approximate way by mapping small volumes in XYZ space to corresponding volumes in CIELAB space and using their bounds instead of exact XYZ errors to drive the splitting and merging of boxes [78, 79].

In Monte Carlo methods, especially when used in scientific simulations of light transport, SPDs are often represented by collections of single-wavelength photons, which are traced individually through a scene in sufficient number for the result to converge to an expected mean [134, 70, 26, 7, 143, 83, 183]. This approach is convenient when rendering dispersive materials with indices of refraction that vary significantly by wavelength. For example, with SPDs represented as n -dimensional vectors of point samples, a faithful rendering of a ray of white light would spawn n single-wavelength rays, or photons, refracted in different directions as it passes through a glass prism. Rendering with photons, however, can be very inefficient, especially in computer graphics contexts, not only because most materials are not dispersive, but also because most of the computation involved in rendering a typical scene is consumed by intersecting rays with geometry, the cost of which is amortized when rays represent spectra instead of photons. Monte Carlo tracing of rays with multiple wavelengths also reduces color noise caused by sampling variance [166, pp. 16-17, 343]. Various strategies have been proposed in the computer graphics literature for stochastic selection of these wavelengths, either when rays are spawned or adaptively (e.g., at dispersive interfaces) [52, 175, 1, 144, 145, 174, 160, 31, 99]. These will not be discussed here, however, as our aim is a representation of spectra that is suitable for rendering in real time, which for some time to come will remain infeasible with Monte Carlo methods.

3.1.2 Polynomials

Motivated by a desire to compute tristimulus coordinates by efficient numerical integration, Moon and Spencer propose a representation of reflectance spectra as degree six polynomials, requiring seven coefficients [124]. This approximates typically smooth reflectance spectra and color signals formed with smooth illuminant SPDs reasonably well, but not color signals formed with fluorescent light sources. Increasing the polynomial degree to capture the narrow peaks in these color signals soon defeats the purpose of dimensionality reduction and makes the method unstable, as the Vandermonde matrix of polynomial basis vectors quickly becomes very ill-conditioned [138].

In a more relevant context, Raso and Fournier [147] propose rendering with reflectance spectra represented as two cubic polynomials, each one requiring four coefficients to approximate half a spectrum. By assuming that a scene contains only a single light source and by computing all reflections *before* illumination is applied to the expanded results, the quality of the approximation is maintained, even when the illuminant contains sharp peaks. Computing a reflection with their method, however, is inefficient, as each of the product’s two degree six polynomials requires 16 multiplications to convolve their coefficients and a 4×6 matrix-vector multiplication to enforce degree closure by projecting them onto the first six Chebyshev polynomials and truncating the higher order terms. Geist et al. [66] improves on this method’s stability and cost by representing both reflectance spectra and illuminant SPDs as linear combinations of the first m Chebyshev polynomials, where m is a power of two, and performing the convolution in the frequency domain using a fast Fourier transform at a reduced cost of $O(m \cdot \log m)$ [38, Ch. 30]. Consistent with results reported by Raso and Fournier, the authors find that a dimensionality of $m \equiv 8$ (degree 7) yields sufficient accuracy when color signals are smooth. They find, however, that rendering with fluorescent lights requires 32 dimensions (degree 31).

3.1.3 Sinusoids

To show that it might be possible to recover reflectance spectra from a small number of sensor responses to corresponding color signals, Wandell estimates the dimensionality of the reflectance spectra of 462 Munsell chips by their projections onto a Fourier basis [171]. He finds that a five-dimensional basis yields a median variance accounted for of 99%. With a similar goal of estimating color signals from a camera’s RGB

response, Drew and Funt project 2,052 color signals, formed with 342 reflectances of natural materials and six daylight SPDs, onto the first three Fourier functions [50]. Although they do not report spectral errors, they find that the approximated color signals yield a mean CIELUV color difference of 11 from camera responses. They also find that the same method yields mean color differences of 33 and 31 when rendering with illuminant A the reflectance spectra of the 24 ColorChecker[®] patches and those of a set of 12 ceramic tiles, respectively. These results, along with Wandell’s, are consistent with observations that although reflectance spectra and natural light source SPDs can be characterized as low-frequency functions [12], the same is less true in general of color signals, especially those formed with sharp-peaked fluorescent light source SPDs [151].

3.1.4 Composite models

Deville et al. describe a method that selects wavelengths to sample by analyzing the SPDs of a scene’s light sources in a pre-process to identify any sharp peaks [46]. In intervals in which all SPDs are smooth, wavelengths are selected to facilitate numerical integration with color matching function by Gaussian quadrature using Meyer’s method. To those wavelengths are added any wavelengths corresponding to spikes, which are integrated separately and added when computing colors. The authors report satisfactory results using six sampled wavelengths, but do not report detailed measurements. Sun et al. also decompose all spectra into smooth backgrounds and collections of spikes [158, 180]. The smooth portions are projected onto a low-order Fourier basis of dimensionality $m - 1$ and resampled at m points into a discrete representation. Rendering an image with a mercury arc lamp light source (an SPD containing six sharp peaks) using $m \equiv 12$ sample points for the smooth spectral backgrounds (yielding 18 sampled wavelengths), the authors report color differences from a ground truth image smaller than those obtained using 64 regularly spaced samples.

3.1.5 Wavelets

Compared to sinusoids, low-dimensional wavelets [114] are well suited for representing low- and high-frequency features of spectra simultaneously [117]. In remote sensing applications they have been used to reduce the dimensionality of large volumes of hyperspectral image data [119, 89, 16, 90, 142, 92, 77]. Claustres et al. use wavelets in spectral rendering, but only for data compression, as reflections are not computed

in the lower-dimensional wavelet basis [35]. That is, they compute reflections by expanding wavelet representations of SPDs and reflectance spectra to their original dimensionality and projecting their products back onto the wavelet basis. Computing products in a wavelet basis is an expensive, nonlinear operation, for which Chern and Wang [30], using a Haar basis, give a formula inspired by an algorithm described by Beylkin [11]. Use of their formula is considerably more expensive than computing products of full-dimensional spectra directly. However, their method renders a sequence of images of increasing accuracy quickly by progressive refinement of approximated spectra. This goal is facilitated by the multi-resolution representation inherent in a wavelet basis.

3.2 Using bases derived empirically

3.2.1 By characteristic vector analysis

Linear models formed with bases derived by an empirical analysis of the data yield, essentially by definition, more accurate approximations than those described in the preceding subsections. That is, for a given choice of $m < n$, such a basis is the answer to the question: In what m -dimensional subspace can the largest projections of our n -dimensional data be found? Being highly correlated at nearby wavelengths, typical spectra have been found to be well approximated in a lower-dimensional subspace [36, 115, 139, 82, 40, 118, 168, 150, 64, 32, 157, 149, 55, 146, 129, 137, 164, 105, 28]. For example, as shown in Fig. 3.1, more than 99% of the variance of Munsell color signals lies within three to five dimensions.

An orthonormal basis for the best-fit subspace, in a least squares sense, is most straightforwardly found by a characteristic vector analysis² of a set of color signals. Let \mathcal{S} be a set of reflectance spectra of p materials to be rendered, \mathcal{E} a set of q light sources SPDs and $\mathcal{C}^{(1)}$ the union of \mathcal{E} and the $p \cdot q$ reflections formed by their Cartesian product:

$$\mathcal{C}^{(1)} \equiv \mathcal{E} \cup (\mathcal{S} \times \mathcal{E}) \equiv \mathcal{E} \cup \{\mathbf{S} \circ \mathbf{E} \mid \mathbf{S} \in \mathcal{S} \text{ and } \mathbf{E} \in \mathcal{E}\}.$$

The singular value decomposition (SVD) of the matrix \mathcal{C} , the columns of which are

²As the term is used here, a *characteristic vector analysis* of a data matrix is equivalent to a principal component analysis (PCA) without first subtracting the mean vector, a procedure sometimes referred to as *uncentered* PCA [14, 106].

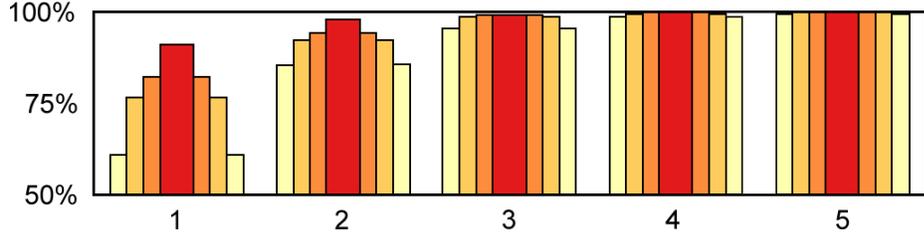


Figure 3.1: Percentages of total variance cumulatively accounted for by the subspaces spanned by the first five characteristic vectors of Munsell color signals. From inner (darker) to outer (lighter) bars, these color signals were formed by reflecting illuminants A, F2, D65 or all three.

the elements of $\mathcal{C}^{(1)}$, is $\mathcal{C} = \mathbf{U}^{(e)} \boldsymbol{\Sigma}^{(e)} \mathbf{V}^{(e)\text{T}}$, where $\mathbf{U}^{(e)}$ and $\mathbf{V}^{(e)}$ are its left and right singular vectors, and $\boldsymbol{\Sigma}^{(e)} \equiv \text{diag}(\boldsymbol{\sigma}^{(e)})$ is the diagonal matrix of its singular values in non-increasing order. By forming an orthonormal basis \mathbf{B} of \mathbb{R}^m with the first m columns of $\mathbf{U}^{(e)}$, a reflection $\mathbf{C} \equiv \mathbf{S} \circ \mathbf{E}$ can be represented as an m -dimensional vector of basis coefficients, $\mathbf{c} \equiv \mathbf{B}^{\text{T}} \mathbf{C}$, and approximated by its projection onto \mathbf{B} as $\mathbf{C} \approx \mathbf{B} \mathbf{B}^{\text{T}} \mathbf{C} = \mathbf{B} \mathbf{c}$, with a total error, $\|\mathbf{C} - \mathbf{B} \mathbf{B}^{\text{T}} \mathbf{C}\|$, that decreases with m . In some cases, an approximated spectrum may contain negative elements of small magnitude. These magnitudes are typically small enough to be simply clamped to 0 without significant consequence in applications requiring physically correct spectra.

This procedure for obtaining \mathbf{B} can be modified by weighting the spectra in a set according to their relative importance or by weighting the whole set according to the relative importance of certain wavelengths [88, c.f. Section 14.2.1]. If, for example, a small number of materials are expected to make a disproportionately large contribution to a scene, their approximations by their projections onto \mathbf{B} can be made more accurate, at the expense of the accuracy of other approximated spectra, by weighting them accordingly before inclusion in \mathcal{C} . Likewise, if certain wavelengths are expected to make a larger contribution to a final image, \mathbf{B} can be obtained instead from $\text{diag}(\mathbf{W}) \mathcal{C}$, where the entries of the weight vector \mathbf{W} are determined by the wavelengths' relative importance. This approach has been used to improve color accuracy attained for a given dimensionality m by constructing \mathbf{W} according to some measure of human visual sensitivity [115, 93, 101, 2, 20]. Because our more general goal, however, is accurate approximation of spectra, not colors, we weight neither spectra nor wavelengths when computing \mathbf{B} .

When rendering with indirect illumination, a recursive computation yields the set $\mathcal{C}^{(k)}$, consisting of the $s \equiv q \sum_{i=0}^k p^i = q(1 - p^{k+1}) / (1 - p)$ SPDs formed by up to k

reflections,

$$\mathcal{C}^{(k)} \equiv \bigcup_{i=0}^k \mathcal{C}^{(i)} = \mathcal{E} \cup \bigcup_{i=1}^k \mathcal{S} \times \mathcal{C}^{(i-1)},$$

the elements of which populate the columns of \mathcal{C} . Because the size of $\mathcal{C}^{(k)}$ increases exponentially with k , when $k > 2$ the directions of the leading left singular values of \mathcal{C} are determined mostly by the highest order reflections, whose approximations by their projections onto $\mathbf{B}^{(k)}$ will be as accurate as possible, at the expense of the accuracy of lower-order reflections. However, in a typical context the contributions to a rendered scene of first and second reflections are of greater importance, especially considering that incident light is usually attenuated as it travels from its source (e.g., by the inverse square or some other function of distance). Moreover, errors incurred on early reflections are propagated to later reflections. For these reasons, when $k > 2$ we compute a $\mathbf{B}^{(2)}$ basis and use it to approximate all higher-order reflections as well. This approach provides a best-fit for first and second reflections and should provide, for typical spectra, a reasonably good fit for further reflections. Thus, in the remainder of this discussion it is assumed that third and subsequent reflections are not included in the SPDs used to construct a basis, which, for convenience, will be denoted simply $\mathbf{B} \equiv \mathbf{B}^{(2)}$, or, if $k \equiv 1$, then $\mathbf{B} \equiv \mathbf{B}^{(1)}$.

Data compression is maximized when both reflectance spectra and SPDs are represented by m -dimensional basis coefficients. Because they span different subspaces (unless $\mathbf{E} \equiv \alpha \mathbf{1}$ for some constant scalar α), accuracy is increased by using a different basis for each set, say $\mathbf{B}^{(s)}$ and $\mathbf{B}^{(e)}$. Computing an m -dimensional reflection with this model, however, is expensive, requiring $6 \cdot m \cdot n - n - m$ floating point operations:

$$\begin{aligned} \hat{\mathbf{c}} &\equiv \mathbf{B}^{(e)\text{T}}[(\mathbf{B}^{(s)}\mathbf{s}) \circ (\mathbf{B}^{(e)}\mathbf{e})] \\ &= \mathbf{B}^{(e)\text{T}}[(\mathbf{B}^{(s)}\mathbf{B}^{(s)\text{T}}\mathbf{S}) \circ (\mathbf{B}^{(e)}\mathbf{B}^{(e)\text{T}}\mathbf{E})] \\ &\approx \mathbf{B}^{(e)\text{T}}(\mathbf{S} \circ \mathbf{E}) = \mathbf{B}^{(e)\text{T}}\mathbf{C} = \mathbf{c}. \end{aligned} \tag{3.1}$$

By instead letting $\mathbf{B} \equiv \mathbf{B}^{(1)}$ or $\mathbf{B}^{(2)}$, as above, and representing reflectance spectra as precomputed $m \times m$ matrices $\{\mathbf{R}^{(i)}\}_{i=1}^p$, this cost can be reduced to $m^2 - m$ operations, at a data storage cost of m^2 for each reflectance. A sequence of k reflections is then

computed with k matrix-vector multiplications:

$$\mathbf{c} \approx [\mathbf{B}^T \text{diag}(\mathcal{S}^{(k)})\mathbf{B}] \cdots [\mathbf{B}^T \text{diag}(\mathcal{S}^{(1)})\mathbf{B}]\mathbf{e} \equiv \mathbf{R}^{(k)} \cdots \mathbf{R}^{(1)}\mathbf{e}.$$

This approach, which is the spectral rendering method described by Peercy [140], is also more accurate than the two-bases approach expressed by (3.1), as it involves one less projection per reflection. That is, the approximation $\mathbf{S} \approx \mathbf{B}^{(s)}\mathbf{B}^{(s)T}\mathbf{S}$ is avoided. Clearly, however, it is practical only for small values of m , since whenever $m^2 > n$ the straightforward representation of spectra as n -dimensional vectors of point samples is more efficient for storage and computation.

3.2.2 By other methods

For completeness it should be mentioned here that other empirical methods have been used to find a lower-dimensional basis to approximate spectra, including independent component analysis [100, 169, 146, 137, 179], nonnegative matrix factorization (NMF) [17, 146, 179, 178] and even neural networks [165, 107, 74, 136, 146, 137]. While there may be theoretical or practical reasons to use such a method (e.g., in the case of NMF, the goal may be to design optical filters for measuring color signals), they are not of interest here, as characteristic vector analysis, a well defined, deterministic and relatively simple technique, is guaranteed to yield a best-fit basis.

Chapter 4

Spectral sharpening

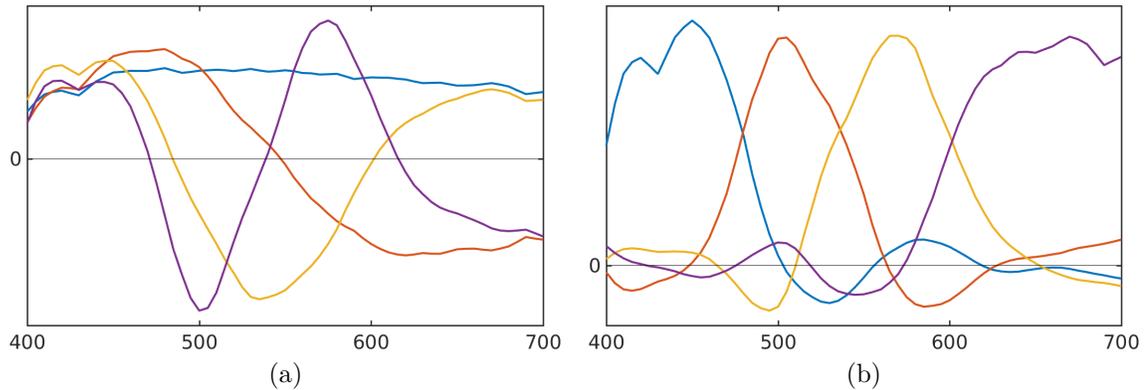


Figure 4.1: (a) Orthonormal basis \mathbf{B} , the first four characteristic vectors of Munsell/D65 $\mathcal{C}^{(1)}$ SPDs. (b) Sharpened basis $\mathbf{Q} \equiv \mathbf{B}\mathbf{T}$.

By transforming an orthonormal basis \mathbf{B} by an $m \times m$ change of basis matrix \mathbf{T} , so that vectors of the “sharp” basis $\mathbf{Q} \equiv \mathbf{B}\mathbf{T}$ are approximately disjoint, or complementary, the method of *spectral sharpening*, as described by Drew and Finlayson [49], reduces the cost of computing a reflection to m multiplications. An example of this transformation is shown in Fig. 4.1. If $\mathbf{Q}_i \circ \mathbf{Q}_j \approx \mathbf{0}$, $i \neq j$, for every pair of transformed basis vectors, then in the new coordinate system the coefficients of a reflection $\mathbf{C} \equiv \mathbf{S} \circ \mathbf{E}$ can be approximated with sufficient accuracy by the product of the coefficients of the reflectance spectrum \mathbf{S} and the incident SPD \mathbf{E} . Because \mathbf{T} is in general not orthogonal, sharp basis coefficients are computed with the Moore-Penrose pseudoinverse $\mathbf{Q}^+ \equiv (\mathbf{Q}^T \mathbf{Q})^{-1} \mathbf{Q}^T = \mathbf{T}^{-1} \mathbf{B}^T$. After transformation back to the standard

basis, the coefficients of the product are expanded to an approximating SPD:

$$\begin{aligned}\tilde{\mathbf{s}} &\equiv \mathbf{Q}^+\mathbf{S} = \mathbf{T}^{-1}\mathbf{s}, & \tilde{\mathbf{e}} &\equiv \mathbf{Q}^+\mathbf{E} = \mathbf{T}^{-1}\mathbf{e}, & \tilde{\mathbf{c}} &\equiv \tilde{\mathbf{s}} \circ \tilde{\mathbf{e}} \\ \mathbf{C} &\approx \mathbf{B}\mathbf{c} \approx \mathbf{B}(\mathbf{T}\tilde{\mathbf{c}}) = \mathbf{Q}\tilde{\mathbf{c}}.\end{aligned}$$

The product of k reflections is computed in the obvious way,

$$\tilde{\mathbf{c}} = \tilde{\mathbf{s}}^{(k)} \circ \dots \circ \tilde{\mathbf{s}}^{(1)} \circ \tilde{\mathbf{e}},$$

and transformed to a linear color by a matrix $\widetilde{\mathbf{M}}$:

$$\mathbf{b} \approx \widetilde{\mathbf{M}}\tilde{\mathbf{c}} \equiv \mathbf{M}\mathbf{Q}\tilde{\mathbf{c}}.$$

Rendering with sharp basis coefficients is much like rendering with RGB triplets or with n -dimensional spectra, but in m dimensions instead of three or n , and with a $3 \times m$ instead of a $3 \times n$ matrix-vector multiplication at each pixel before display.

In [49] Drew and Finlayson reverse the roles of \mathbf{Q} and $\mathbf{Q}^{+\text{T}}$. That is, coefficients are computed and expanded as

$$\mathbf{C} \approx \mathbf{Q}^{+\text{T}}\tilde{\mathbf{c}} = \mathbf{Q}^{+\text{T}}(\tilde{\mathbf{s}} \circ \tilde{\mathbf{e}}) = \mathbf{Q}^{+\text{T}}[(\mathbf{Q}^{\text{T}}\mathbf{S}) \circ (\mathbf{Q}^{\text{T}}\mathbf{E})].$$

Since $(\mathbf{Q}^{+\text{T}})^+ = \mathbf{T}^{\text{T}}\mathbf{B}^{\text{T}} = \mathbf{Q}^{\text{T}}$, the two formulations are mathematically equivalent; the role of our matrix \mathbf{T} is equivalent to that of their $\mathbf{T}^{-\text{T}}$. We choose ours because it is standard and intuitive. Algorithmically, however, there is an important difference, because Drew and Finlayson seek, in our formulation, a matrix $\mathbf{T}^{-\text{T}}$ to sharpen $\mathbf{Q}^{+\text{T}}$. Results obtained by each method will be different but similar, because sharpening $\mathbf{Q}^{+\text{T}}$ tends to sharpen \mathbf{Q} indirectly, and vice versa. This can be understood by considering that if $\mathbf{Q}_i \circ \mathbf{Q}_j \approx \mathbf{0}$, then $\langle \mathbf{Q}_i, \mathbf{Q}_j \rangle \approx 0$, and

$$\mathbf{Q}^{\text{T}}\mathbf{Q} = \mathbf{T}^{\text{T}}\mathbf{B}^{\text{T}}\mathbf{B}\mathbf{T} = \mathbf{T}^{\text{T}}\mathbf{T} \approx \text{diag}(\mathbf{d})$$

for some vector $\mathbf{d} \equiv (d_1, d_2, \dots, d_m)^{\text{T}}$. Thus,

$$\mathbf{Q}_i^{+\text{T}} \circ \mathbf{Q}_j^{+\text{T}} \approx (d_i d_j)^{-1} \mathbf{Q}_i \circ \mathbf{Q}_j \approx \mathbf{0}.$$

Results reported in [49], obtained by sharpening $\mathbf{Q}^{+\text{T}}$, are thus due to indirect sharpening of \mathbf{Q} , which we sharpen directly. Another difference between our version of

spectral sharpening and that of Drew and Finlayson is that they constrain \mathbf{Q}^+ to be nonnegative, while we do not. The authors offer no justification for this constraint, other than the claim that it “makes intuitive sense.” Not only is it mathematically unnecessary, but the freedom afforded by its removal can only improve the approximations computed with a sharp basis.

A related method described by Darling also computes reflections by multiplying sharp basis coefficients [43, 41]. Darling’s ostensibly universal basis, intended for general use in real-time rendering with any spectra, is constructed by assuming that the rows of \mathbf{Q}^+ consist of approximately disjoint Gaussian functions, the peak wavelengths and widths of which are found by minimizing the color differences of the approximations of a set of $\mathcal{C}^{(1)}$ reflections constructed from 100 reflectance spectra of artist’s paints and 19 various standard illuminants. The dimensionality of the basis is fixed at $m \equiv 6$ to accommodate standard rendering software by allowing coefficients to be referenced as two sets of RGB triplets. As is shown in the following chapters, by making no assumptions regarding the form of \mathbf{Q}^+ or \mathbf{Q} , nor therefore that of \mathbf{T} , and by optimizing \mathbf{Q} for a particular set of spectra to be rendered, including secondary reflections, we can expect better approximations, especially if the dimensionality of the basis is allowed to vary to attain a desired accuracy.

Scaling sharp basis coefficients relative to a maximum reflectance of $\mathbf{S} \equiv \mathbf{1}$ by the normalization $\mathbf{Q}^+\mathbf{1} = \mathbf{1}$, or, equivalently, $\mathbf{T}\mathbf{1} = \mathbf{B}^T\mathbf{1}$, ensures that reflection by a white surface, or illumination by an equal-energy light source, incurs no additional error, as the n -dimensional unit is projected to the unit in m dimensions:

$$\begin{aligned}\mathbf{Q}[(\mathbf{Q}^+\mathbf{1}) \circ (\mathbf{Q}^+\mathbf{E})] &= \mathbf{Q}(\mathbf{1} \circ \tilde{\mathbf{e}}) = \mathbf{Q}\tilde{\mathbf{e}} = (\mathbf{B}\mathbf{T})(\mathbf{T}^{-1}\mathbf{B}^T\mathbf{E}) = \mathbf{B}\mathbf{B}^T\mathbf{E} \\ \mathbf{Q}[(\mathbf{Q}^+\mathbf{S}) \circ (\mathbf{Q}^+\mathbf{1})] &= \mathbf{Q}(\tilde{\mathbf{s}} \circ \mathbf{1}) = \mathbf{Q}\tilde{\mathbf{s}} = (\mathbf{B}\mathbf{T})(\mathbf{T}^{-1}\mathbf{B}^T\mathbf{S}) = \mathbf{B}\mathbf{B}^T\mathbf{S}.\end{aligned}$$

In the constrained optimization problems described in the following chapter, this constraint is imposed explicitly by normalizing a solution. Solutions to the unconstrained problems, on the other hand, are typically found to be similarly constrained, or to a close approximation. That this should be the case will be better understood after the discussion that follows.

Chapter 5

Finding the change of basis matrix

To find \mathbf{T} an optimization problem must be posed and solved. The most obvious candidate is the minimization of a measure of the total error incurred when a set of reflections $\mathcal{C}^{(k)}$ is approximated by multiplying sharp basis coefficients:

$$\mathbf{T}^{(1)} \equiv \arg \min_{\mathbf{T}} f^{(1)}(\mathbf{T}, \mathbf{B}, \mathcal{S}, \mathcal{C}), \quad (5.1)$$

where $\mathcal{C} \equiv \mathcal{C}^{(1)}$, $\mathbf{c} \equiv \mathbf{B}^T \mathcal{C}$, $\tilde{\mathbf{c}} \equiv \mathbf{T}^{-1} \mathbf{c}$, $\tilde{\mathcal{J}} \equiv \mathbf{T}^{-1} \mathcal{J}$ and

$$f^{(1)} \equiv \sum_{i=1}^p \sum_{j=1}^{q \frac{1-p^k}{1-p}} \|\mathbf{B} \mathbf{B}^T \text{diag}(\mathcal{S}_i) \mathbf{B} \mathbf{B}^T \mathcal{C}_j - \mathbf{Q}(\tilde{\mathcal{J}}_i \circ \tilde{\mathbf{c}}_j)\|_F^2 \quad (5.2)$$

$$= \sum_{i=1}^p \sum_{j=1}^{q \frac{1-p^k}{1-p}} \|\mathbf{B}^T \text{diag}(\mathcal{S}_i) \mathbf{B} - \mathbf{T} \text{diag}(\tilde{\mathcal{J}}_i) \mathbf{T}^{-1}\|_F^2 \mathbf{c}_j^2 \quad (5.3)$$

$$= \sum_{i=1}^p \|\mathbf{B}^T \text{diag}(\mathcal{S}_i) \mathbf{B} - \mathbf{T} \text{diag}(\tilde{\mathcal{J}}_i) \mathbf{T}^{-1}\|_F^2, \quad (5.4)$$

where $\|\cdot\|_F$ is the Frobenius norm, which is invariant under orthogonal transformation, hence (5.3), allowing the error to be minimized in m instead of n dimensions. The columns of the $m \times q(1-p^k)/(1-p)$ matrix $\mathbf{c} \equiv \mathbf{c}^{(k-1)}$ contain basis \mathbf{B} coefficients of SPDs of light sources and possibly first reflections, depending on whether the illumination model is local ($k \equiv 1$) or global ($k > 1$). In the latter case we set $k \equiv 2$ to optimize for first and second reflections. This is justified by the smaller contribution to a rendered scene made by further reflections, which tend toward small-norm constant functions. Because their number increases exponentially with k , they would overwhelmingly influence solutions to (5.1) at the expense of first and

second reflections, which are of greater importance, as noted in the preceding chapter.

This is an unconstrained, nonlinear least squares problem, involving m^2 parameters, the entries of \mathbf{T} , and $p \cdot m \cdot s$ squared residuals, where $s \equiv q \cdot (1 - p^k)/(1 - p)$, each computed at an amortized cost of $O(m)$. When \mathcal{S} is the 1,269 Munsell reflectance spectra, $\mathcal{E} \equiv \{\text{D65, A, F2}\}$, $k \equiv 2$ and $m \equiv 4$, this amounts to 19,339,560 residuals. In practice, this number would be much larger, as a typical scene would involve many more materials. When represented with the same dimensionality, one million reflectance spectra (a not unreasonably large number) and the same number of illuminants would involve an unwieldy 1.2×10^{13} residuals.

This number can be dramatically reduced by posing a much smaller equivalent problem suggested by two observations. First, noting that $\|\mathbf{X}\|_F^2 = \text{tr}(\mathbf{X}^T \mathbf{X}) = \text{tr}(\mathbf{X} \mathbf{X}^T)$, an equivalent objective function, involving $p \cdot m^2$ squared residuals, is

$$f^{(1)} = \sum_{i=1}^p \|\mathbf{B}^T \text{diag}(\mathcal{S}_i) \mathbf{B} - \mathbf{T} \text{diag}(\tilde{\mathcal{S}}_i) \mathbf{T}^{-1}\| \sqrt{\mathbf{c} \mathbf{c}^T} \|_F^2, \quad (5.5)$$

where $\sqrt{\mathbf{c} \mathbf{c}^T}$ is well defined, as $\mathbf{c} \mathbf{c}^T$ is positive-semidefinite. Second, each of the p terms in (5.5) can be rearranged as the norm of a linear transformation of \mathcal{S}_i :

$$f^{(1)} = \sum_{i=1}^p \|\{(\mathbf{B} \sqrt{\mathbf{c} \mathbf{c}^T})^T * \mathbf{B}^T - [(\mathbf{T}^{-1} \sqrt{\mathbf{c} \mathbf{c}^T})^T * \mathbf{T}] \mathbf{Q}^+\} \mathcal{S}_i\|_F^2 \equiv \sum_{i=1}^p \|\mathbf{D} \mathcal{S}_i\|_F^2. \quad (5.6)$$

The $m^2 \times n$ matrix $(\mathbf{B} \sqrt{\mathbf{c} \mathbf{c}^T})^T * \mathbf{B}^T$ and $m^2 \times m$ matrix $(\mathbf{T}^{-1} \sqrt{\mathbf{c} \mathbf{c}^T})^T * \mathbf{T}$ in (5.6) are Khatri-Rao products, employed here for notational convenience. The Khatri-Rao (or columnwise Kronecker) product of $m \times n$ matrix \mathbf{X} and $p \times n$ matrix \mathbf{Y} is the $m \cdot p \times n$ matrix:

$$\mathbf{X} * \mathbf{Y} \equiv [(\mathbf{X}_1 \otimes \mathbf{Y}_1) \quad (\mathbf{X}_2 \otimes \mathbf{Y}_2) \quad \cdots \quad (\mathbf{X}_n \otimes \mathbf{Y}_n)],$$

where \otimes denotes the Kronecker product:

$$\mathbf{X}_i \otimes \mathbf{Y}_i \equiv \begin{bmatrix} X_{1i} \\ \vdots \\ X_{mi} \end{bmatrix} \otimes \mathbf{Y}_i \equiv \begin{bmatrix} X_{1i} \mathbf{Y}_i \\ \vdots \\ X_{mi} \mathbf{Y}_i \end{bmatrix}.$$

For example, if

$$\mathbf{X} \equiv \begin{bmatrix} X_{11} & X_{12} & \cdots & X_{1n} \\ \vdots & \vdots & \ddots & \vdots \\ X_{m1} & X_{m2} & \cdots & X_{mn} \end{bmatrix} \quad \text{and} \quad \mathbf{Y} \equiv \begin{bmatrix} Y_{11} & Y_{12} & \cdots & Y_{1n} \\ \vdots & \vdots & \ddots & \vdots \\ Y_{p1} & Y_{p2} & \cdots & Y_{pn} \end{bmatrix}$$

then

$$\mathbf{X} * \mathbf{Y} \equiv \begin{bmatrix} X_{11}Y_{11} & X_{12}Y_{12} & \cdots & X_{1n}Y_{1n} \\ \vdots & \vdots & \ddots & \vdots \\ X_{11}Y_{p1} & X_{12}Y_{p2} & \cdots & X_{1n}Y_{pn} \\ X_{21}Y_{21} & X_{22}Y_{22} & \cdots & X_{2n}Y_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ X_{21}Y_{p1} & X_{22}Y_{p2} & \cdots & X_{2n}Y_{pn} \\ \vdots & \vdots & \ddots & \vdots \\ X_{m1}Y_{11} & X_{m2}Y_{12} & \cdots & X_{mn}Y_{1n} \\ \vdots & \vdots & \ddots & \vdots \\ X_{m1}Y_{p1} & X_{m2}Y_{p2} & \cdots & X_{mn}Y_{pn} \end{bmatrix} = \begin{bmatrix} \mathbf{X}_1^T \circ \mathbf{Y}_1^T \\ \vdots \\ \mathbf{X}_1^T \circ \mathbf{Y}_p^T \\ \mathbf{X}_2^T \circ \mathbf{Y}_1^T \\ \vdots \\ \mathbf{X}_2^T \circ \mathbf{Y}_p^T \\ \vdots \\ \mathbf{X}_n^T \circ \mathbf{Y}_1^T \\ \vdots \\ \mathbf{X}_n^T \circ \mathbf{Y}_p^T \end{bmatrix}.$$

Thus, since

$$f^{(1)} = \sum_{i=1}^p \|\mathbf{D}\mathcal{S}_i\|_F^2 = \|\mathbf{D}\mathcal{S}\|_F^2 = \text{tr}(\mathbf{D}\mathcal{S}\mathcal{S}^T\mathbf{D}^T) = \text{tr}(\mathbf{D}\mathbf{U}^{(\mathcal{S})}\boldsymbol{\Sigma}^{(\mathcal{S})2}\mathbf{U}^{(\mathcal{S})T}\mathbf{D}^T), \quad (5.7)$$

the sum of p terms in (5.5) is equivalent to the sum of n terms, the first factors of which are linear combinations of the left singular vectors of \mathcal{S} weighted by their respective singular values, instead of linear combinations of the spectra \mathcal{S}_i :

$$f^{(1)} = \sum_{i=1}^n \|\mathbf{B}^T \text{diag}(\mathbf{U}_i^{(\mathcal{S})} \boldsymbol{\sigma}_i^{(\mathcal{S})}) \mathbf{B} - \mathbf{T} \text{diag}(\mathbf{Q}^+ \mathbf{U}_i^{(\mathcal{S})} \boldsymbol{\sigma}_i^{(\mathcal{S})}) \mathbf{T}^{-1}] \sqrt{\mathbf{c}\mathbf{c}^T}\|_F^2 \quad (5.8)$$

$$\equiv \sum_{i=1}^n \|\mathbf{B}^T \text{diag}(\widehat{\mathbf{U}}_i^{(\mathcal{S})}) \mathbf{B} - \mathbf{T} \text{diag}(\widetilde{\mathbf{u}}_i^{(\mathcal{S})}) \mathbf{T}^{-1}] \sqrt{\mathbf{c}\mathbf{c}^T}\|_F^2 \quad (5.9)$$

$$\equiv \sum_{i=1}^n \|\boldsymbol{\Delta}^{(i)} \sqrt{\mathbf{c}\mathbf{c}^T}\|_F^2 = \|\boldsymbol{\Delta} \sqrt{\mathbf{c}\mathbf{c}^T}\|_F^2, \quad (5.10)$$

where the $n \cdot m \times m$ matrix $\boldsymbol{\Delta}$ consists of the n vertical blocks $\boldsymbol{\Delta}^{(i)}$. This reformulation

reduces the size of the problem to $n \cdot m^2$ squared residuals (each computed at an amortized cost of $O(m)$), a number independent of the sizes of \mathcal{S} and \mathcal{E} . The problem's size can be reduced even further to $d \cdot m^2$, where $d \ll n$, by including only the leading d left singular vectors, since the corresponding singular values descend rapidly towards zero. In practice, $8 \lesssim d \lesssim 20$ should suffice for typical reflectance spectra. Thus, for $4 \leq m \leq 9$ and any number of spectra, the problem size can be reduced to 128 to 1,620 squared residuals.

An examination of (5.9) and (5.10) suggests several more quickly solved subproblems whose solutions can approximate $\mathbf{T}^{(1)}$ at the cost of some possibly tolerable reduction in accuracy. In this chapter we identify three of these and discuss methods for solving them as well as $f^{(1)}$. All of these problems, including three methods of explicit spectral sharpening, are summarized in Table 5.1. Results obtained using their solutions are presented in Chapter 6. First, however, in the next section we pose a different problem whose solution yields a coordinate system in which the error incurred by multiplying sharp basis coefficients vanishes when approximating $\mathcal{C}^{(1)}$ reflections of a single light source SPD.

5.1 For $q \equiv k \equiv 1$, by diagonalizing $\mathbf{B}^T \text{diag}(\mathbf{E})^{-1} \mathbf{B}$

When a scene is rendered with only first reflections and all sources of light share the same SPD, every reflectance can be retrieved from the corresponding reflected SPD by factoring out the illumination¹: $\mathbf{S} = \text{diag}(\mathbf{E})^{-1} \mathbf{C}$. In m dimensions a similar operation can be applied to a reflected SPD's coefficients²: $\mathbf{s} \approx \mathbf{B}^T \text{diag}(\mathbf{E})^{-1} \mathbf{B} \mathbf{c}$. Because this operator is symmetric, it can be diagonalized by a matrix \mathbf{V} whose columns are eigenvectors forming an orthonormal basis for \mathbb{R}^m :

$$\mathbf{V}^T \mathbf{B}^T \text{diag}(\mathbf{E})^{-1} \mathbf{B} \mathbf{V} = \text{diag}(\boldsymbol{\mu}),$$

where the entries of $\boldsymbol{\mu}$ are the corresponding eigenvalues. The normalization

$$\mathbf{T}^{(\mathbf{V})} \equiv \mathbf{V} \text{diag}(\mathbf{V}^T \mathbf{e}) \text{diag}(\boldsymbol{\mu})$$

¹Because the typical light source SPD is non-zero at every wavelength, $\text{diag}(\mathbf{E})^{-1}$ exists because $\text{diag}(\mathbf{E})$ is nonsingular. Whenever this is not the case wavelengths where the SPD vanishes make no contribution to any color signal, and may be omitted when rendering, effectively reducing the dimensionality of spectra from n to $n' < n$.

²Throughout this section, in which only $\mathcal{C}^{(1)}$ SPDs are considered, $\mathbf{B} \equiv \mathbf{B}^{(1)}$.

yields a coordinate system in which the illuminant's coordinates are the inverted eigenvalues³:

$$\tilde{\mathbf{e}} \equiv \mathbf{T}^{(\mathbf{V})^{-1}} \mathbf{e} = \text{diag}(\boldsymbol{\mu})^{-1} \text{diag}(\mathbf{V}^T \mathbf{e})^{-1} \mathbf{V}^T \mathbf{e} = \text{diag}(\boldsymbol{\mu})^{-1}.$$

A reflectance spectrum's coordinates are obtained by factoring out the illuminant's,

$$\tilde{\mathbf{s}} \equiv \text{diag}(\tilde{\mathbf{e}})^{-1} \tilde{\mathbf{c}} = \text{diag}(\boldsymbol{\mu}) \mathbf{T}^{(\mathbf{V})^{-1}} \mathbf{c} = \text{diag}(\mathbf{V}^T \mathbf{e})^{-1} \mathbf{V}^T \mathbf{c}.$$

The coordinates of a reflected SPD are thus $\tilde{\mathbf{c}} \equiv \tilde{\mathbf{e}} \circ \tilde{\mathbf{s}}$, and the error incurred when this product is transformed to the standard basis is

$$\begin{aligned} \mathbf{r} &\equiv \mathbf{B}^T \mathbf{C} - \mathbf{T}^{(\mathbf{V})} \tilde{\mathbf{c}} = \mathbf{B}^T \mathbf{C} - \mathbf{V} \text{diag}(\boldsymbol{\mu})^{-1} \mathbf{V}^T \mathbf{B}^T \mathbf{S} \\ &= \mathbf{B}^T \mathbf{C} - [\mathbf{B}^T \text{diag}(\mathbf{E})^{-1} \mathbf{B}]^{-1} \mathbf{B}^T \text{diag}(\mathbf{E})^{-1} \mathbf{C}. \end{aligned}$$

Noting that

$$\begin{aligned} \mathbf{r}' &\equiv \mathbf{B}^T \mathbf{C} - [\mathbf{B}^T \text{diag}(\mathbf{E})^{-1} \mathbf{B}]^{-1} \mathbf{B}^T \text{diag}(\mathbf{E})^{-1} \mathbf{B} \mathbf{B}^T \mathbf{C} \\ &\equiv \mathbf{B}^T \mathbf{C} - [\mathbf{B}^T \text{diag}(\mathbf{E})^{-1} \mathbf{B}]^{-1} \mathbf{s}' = \mathbf{0}, \end{aligned}$$

the approximation can be made exact by rendering instead with coefficients $\tilde{\mathbf{s}}' \equiv \mathbf{Q}^+ \mathbf{S}'$ of transformed reflectance spectra $\mathbf{S}' \equiv \text{diag}(\mathbf{E})^{-1} \mathbf{B} \mathbf{B}^T \mathbf{C}$, formed by factoring the illuminant out of the corresponding SPD's projection onto \mathbf{B} .

In Figs. 5.1-5.18 are shown the first nine vectors of $\mathbf{B}^{(1)}$ for Munsell, Natural Colors and METACOW single-illuminant $\mathcal{C}^{(1)}$ SPDs formed with $\mathbf{E} \in \{\text{D65}, \text{A}, \text{F2}\}$, along with, for dimensionalities $m \equiv 4, \dots, 9$, the corresponding sharp bases $\mathbf{Q}^{(\mathbf{V})} \equiv \mathbf{B} \mathbf{T}^{(\mathbf{V})}$ found by diagonalizing $\mathbf{B}^T \text{diag}(\mathbf{E})^{-1} \mathbf{B}$. Also included, for comparison, are the corresponding sharp bases $\mathbf{Q}^{(5)} \equiv \mathbf{B} \mathbf{T}^{(5)}$ found by spectral sharpening (see Section 5.4). As shown by the corresponding values of $f^{(5)}$ measuring sharpness, in the rightmost columns of plots, the sharp bases found by the method described in this section are considerably less sharp, although, as will be shown in the next chapter, they yield lower approximation errors.

When rendering with indirect illumination, or with more than one type of light, every reflected SPD cannot be factored as the product of a single light source SPD and a reflectance spectrum. In these cases \mathbf{T} is found by minimizing $f^{(1)}$, the total

³Because $\mathbf{B}^T \text{diag}(\mathbf{E})^{-1} \mathbf{B}$ is positive definite, its eigenvalues are positive and $\text{diag}(\boldsymbol{\mu})^{-1}$ exists.

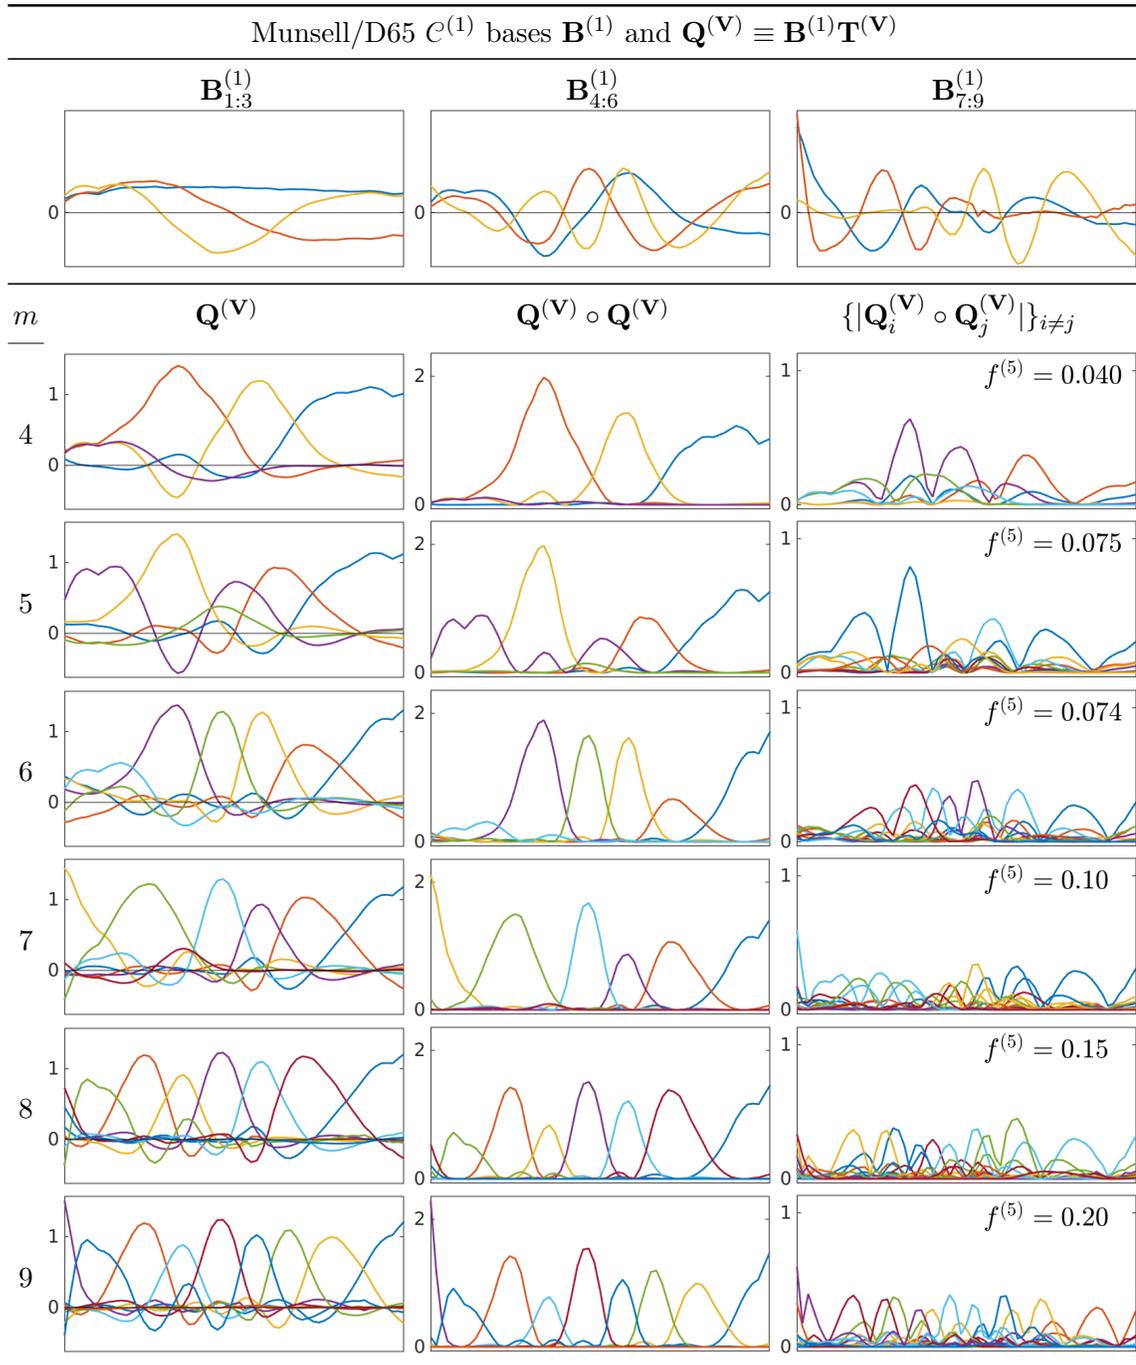


Figure 5.1: For Munsell $\mathcal{C}^{(1)}$ color signals formed with D65, the orthonormal basis $\mathbf{B}^{(1)}$, the sharp bases $\mathbf{Q}^{(\mathbf{V})} \equiv \mathbf{B}^{(1)}\mathbf{T}^{(\mathbf{V})}$ (for dimensionalities $m \equiv 4, \dots, 9$), their squares and the absolute values of the pairwise products of their columns.

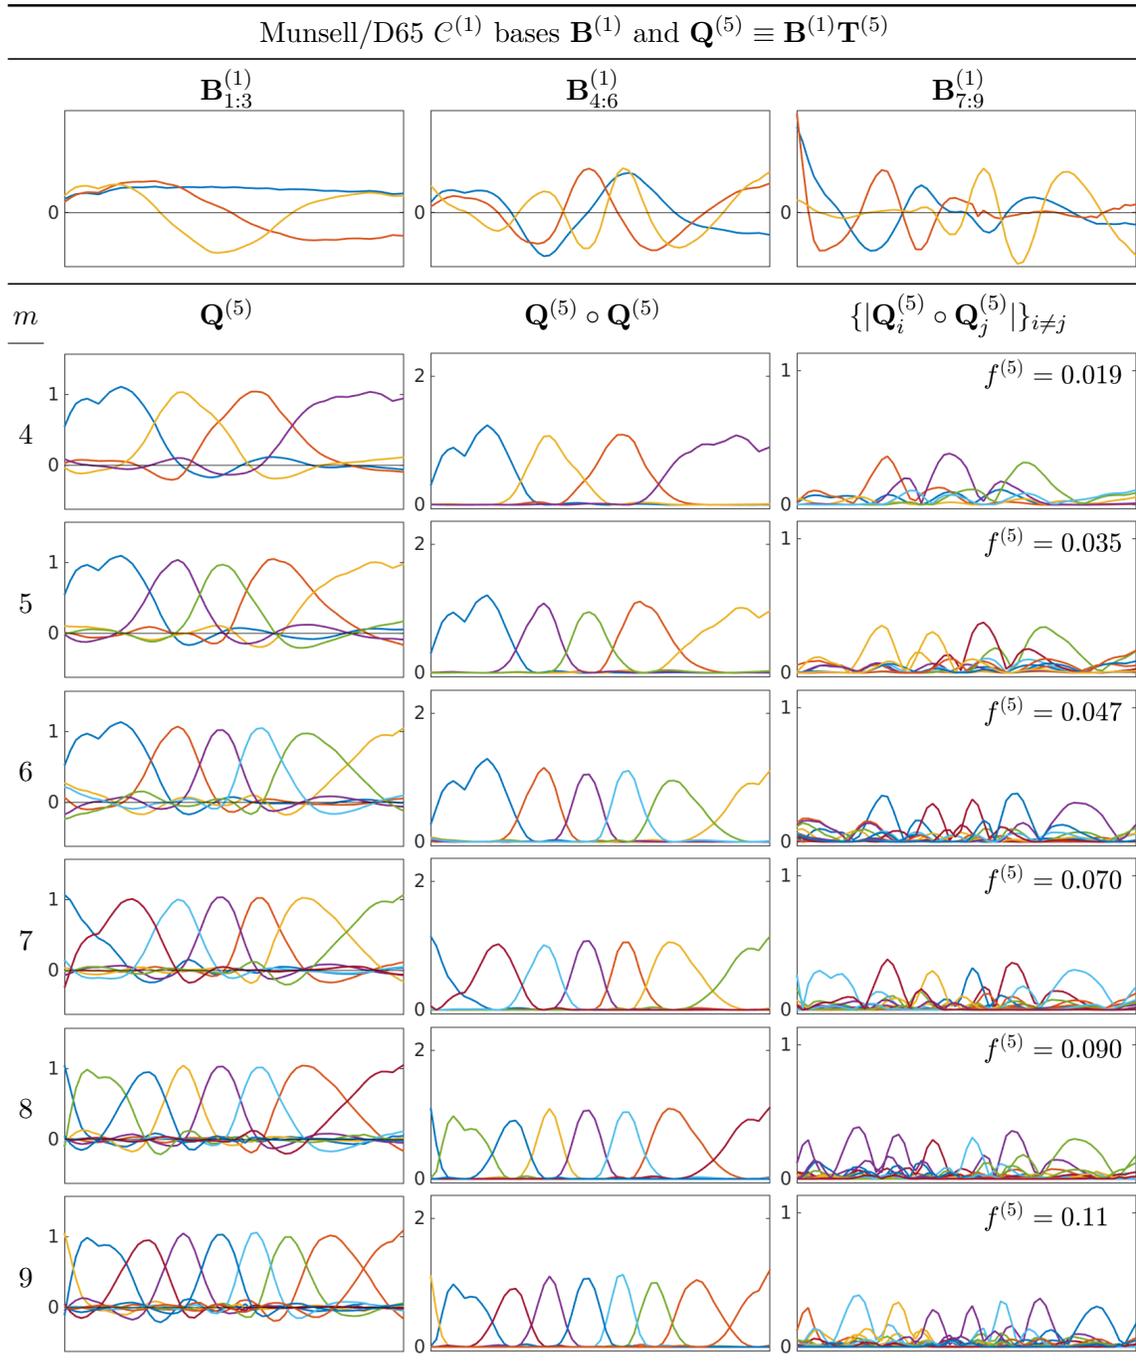


Figure 5.2: For Munsell $\mathcal{C}^{(1)}$ color signals formed with D65, the orthonormal basis $\mathbf{B}^{(1)}$, the sharp bases $\mathbf{Q}^{(5)} \equiv \mathbf{B}^{(1)}\mathbf{T}^{(5)}$ (for dimensionalities $m \equiv 4, \dots, 9$), their squares and the absolute values of the pairwise products of their columns.

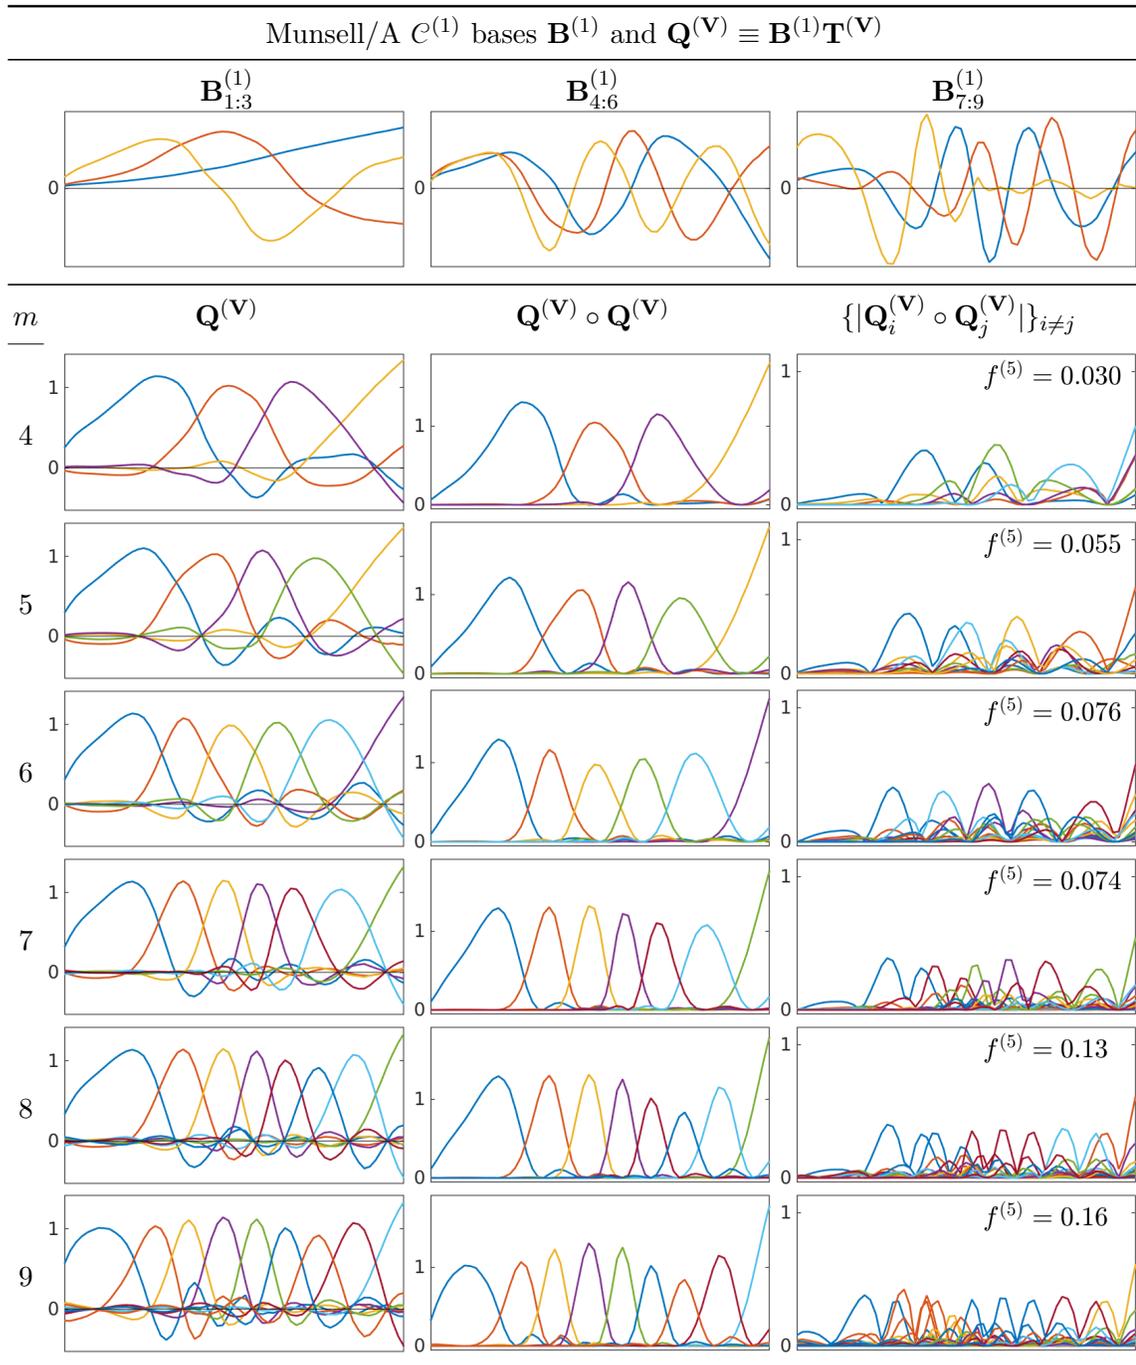


Figure 5.3: For Munsell $\mathcal{C}^{(1)}$ color signals formed with illuminant A, the orthonormal basis $\mathbf{B}^{(1)}$, the sharp bases $\mathbf{Q}^{(\mathbf{V})} \equiv \mathbf{B}^{(1)}\mathbf{T}^{(\mathbf{V})}$ (for dimensionalities $m \equiv 4, \dots, 9$), their squares and the absolute values of the pairwise products of their columns.

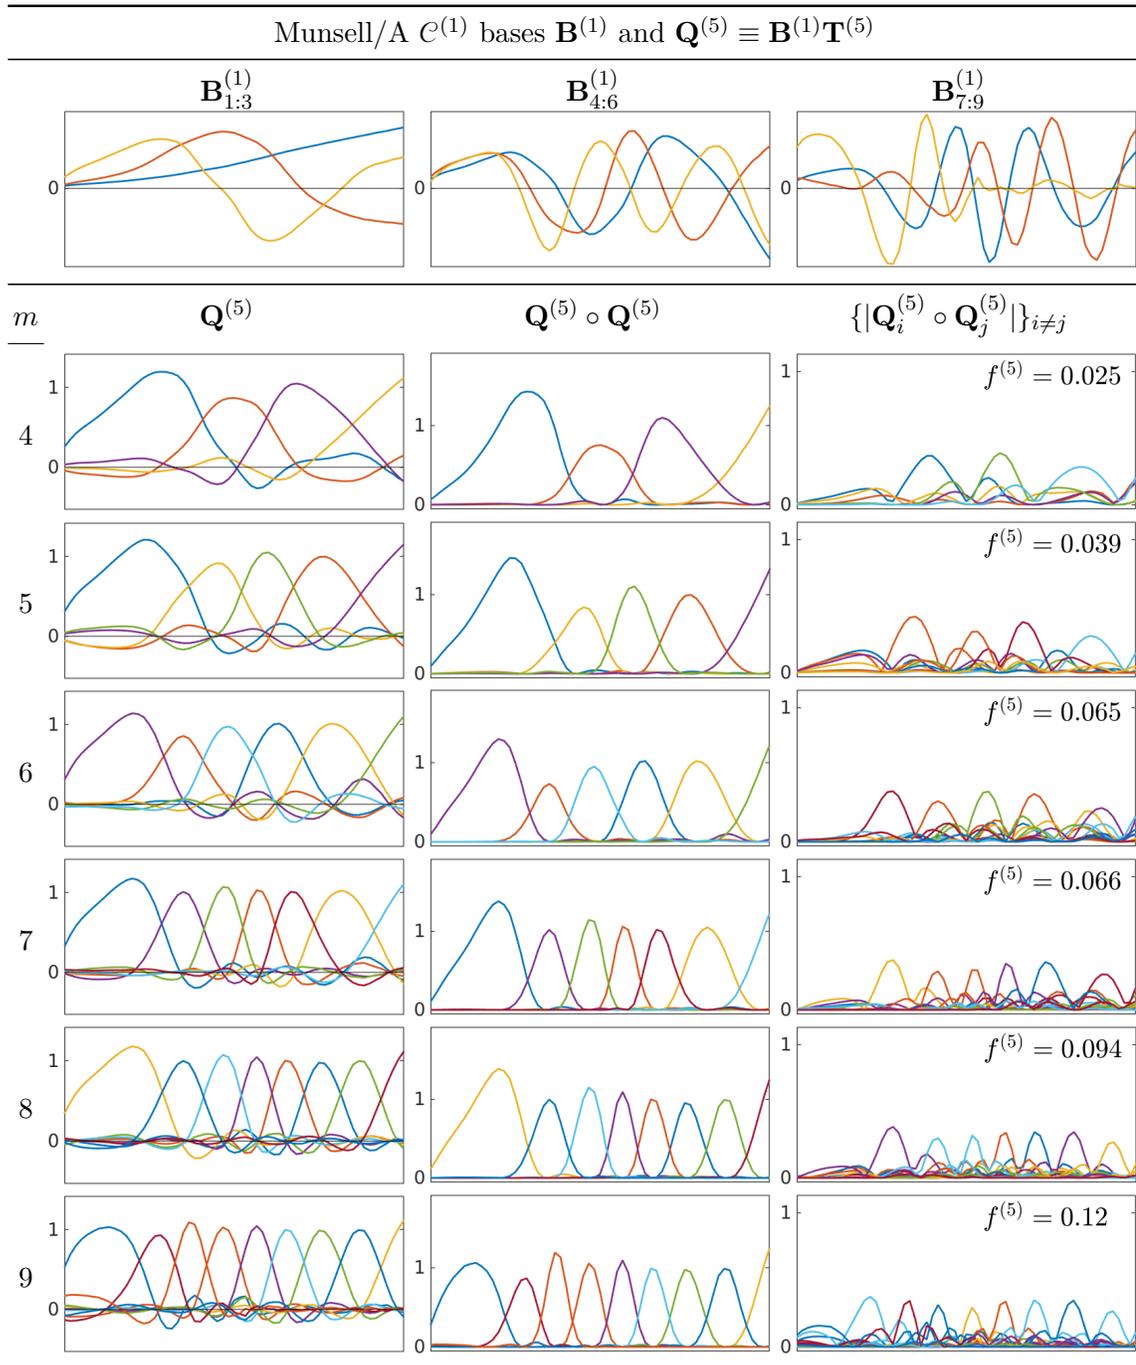


Figure 5.4: For Munsell $\mathcal{C}^{(1)}$ color signals formed with illuminant A, the orthonormal basis $\mathbf{B}^{(1)}$, the sharp bases $\mathbf{Q}^{(5)} \equiv \mathbf{B}^{(1)}\mathbf{T}^{(5)}$ (for dimensionalities $m \equiv 4, \dots, 9$), their squares and the absolute values of the pairwise products of their columns.

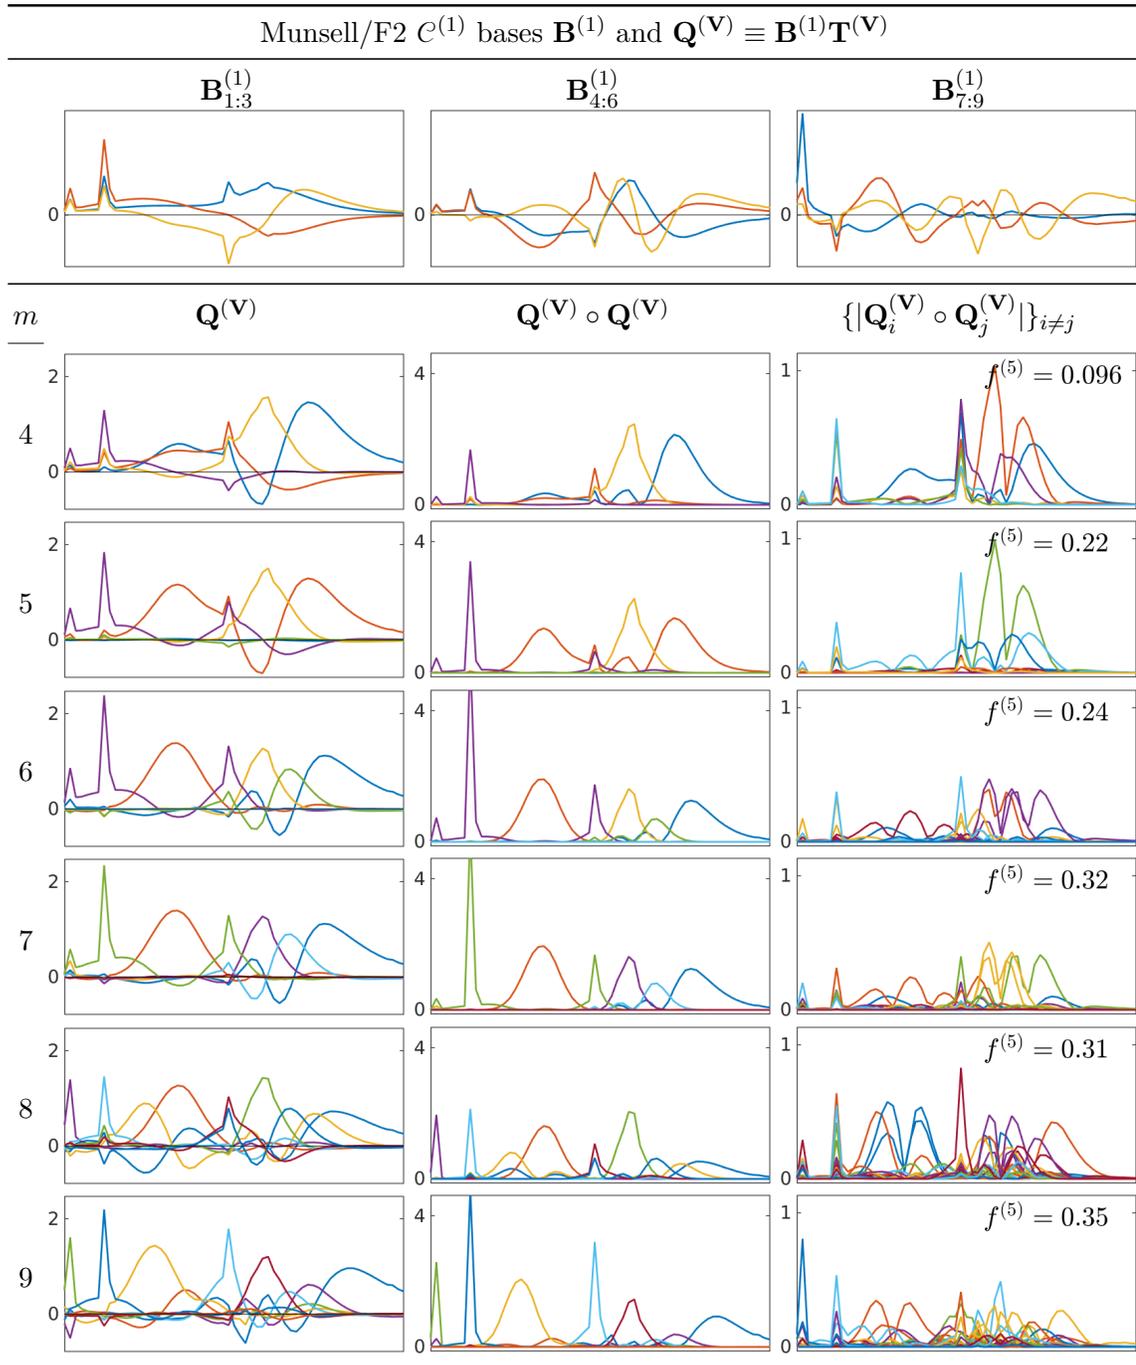


Figure 5.5: For Munsell $\mathcal{C}^{(1)}$ color signals formed with F2, the orthonormal basis $\mathbf{B}^{(1)}$, the sharp bases $\mathbf{Q}^{(\mathbf{V})} \equiv \mathbf{B}^{(1)}\mathbf{T}^{(\mathbf{V})}$ (for dimensionalities $m \equiv 4, \dots, 9$), their squares and the absolute values of the pairwise products of their columns.

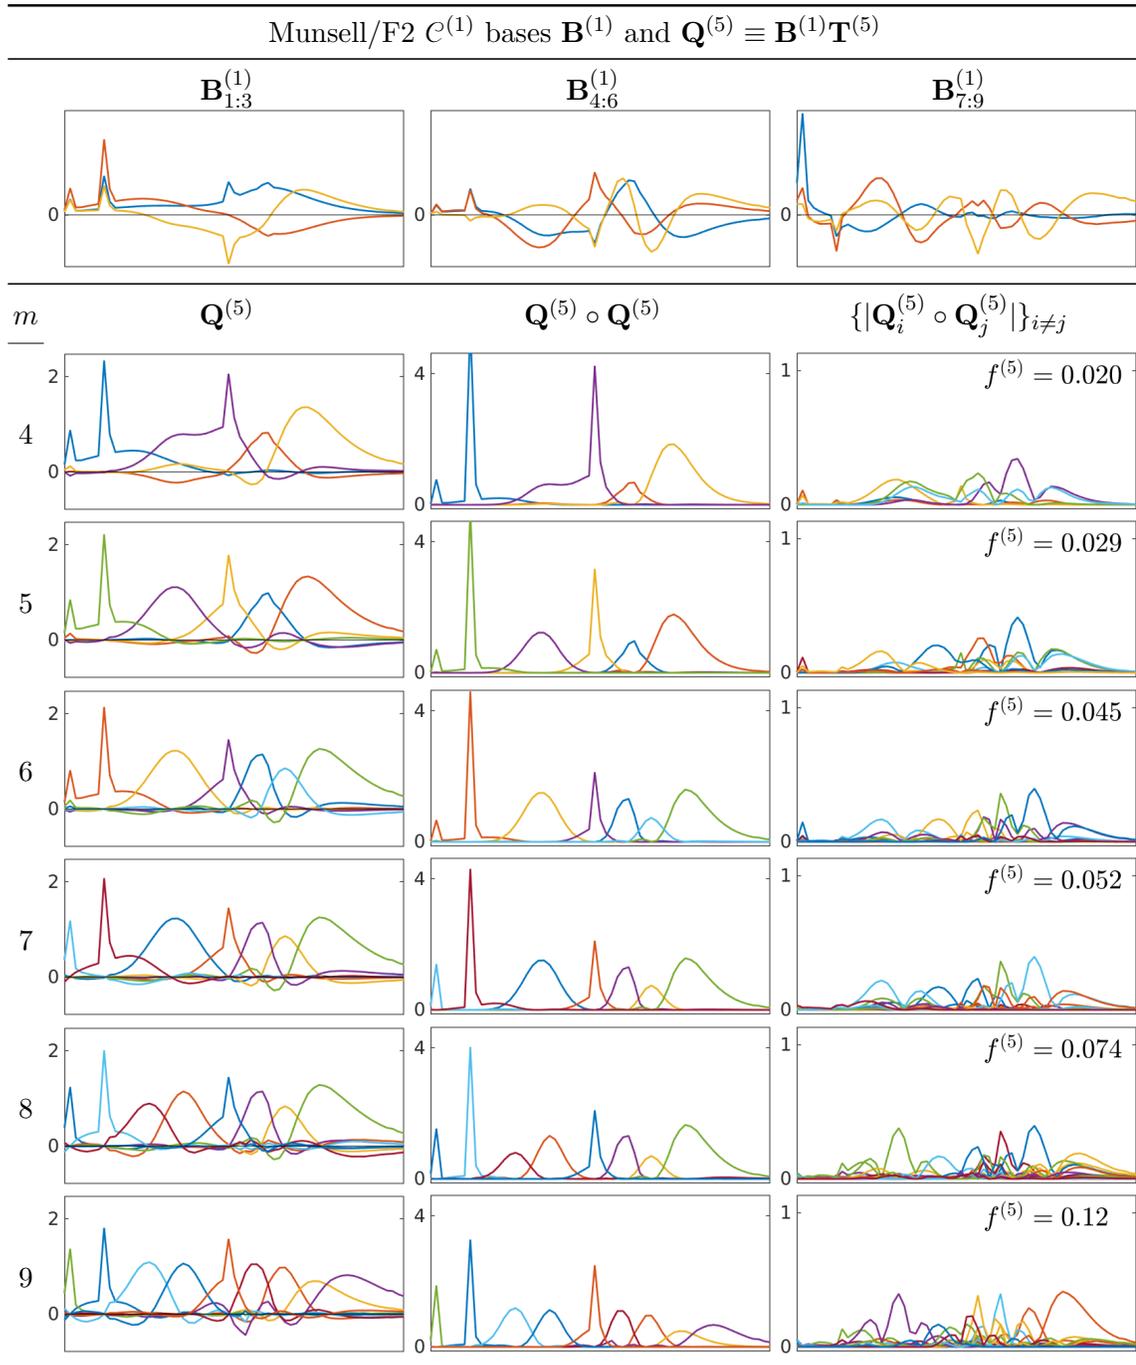


Figure 5.6: For Munsell $\mathcal{C}^{(1)}$ color signals formed with F2, the orthonormal basis $\mathbf{B}^{(1)}$, the sharp bases $\mathbf{Q}^{(5)} \equiv \mathbf{B}^{(1)}\mathbf{T}^{(5)}$ (for dimensionalities $m \equiv 4, \dots, 9$), their squares and the absolute values of the pairwise products of their columns.

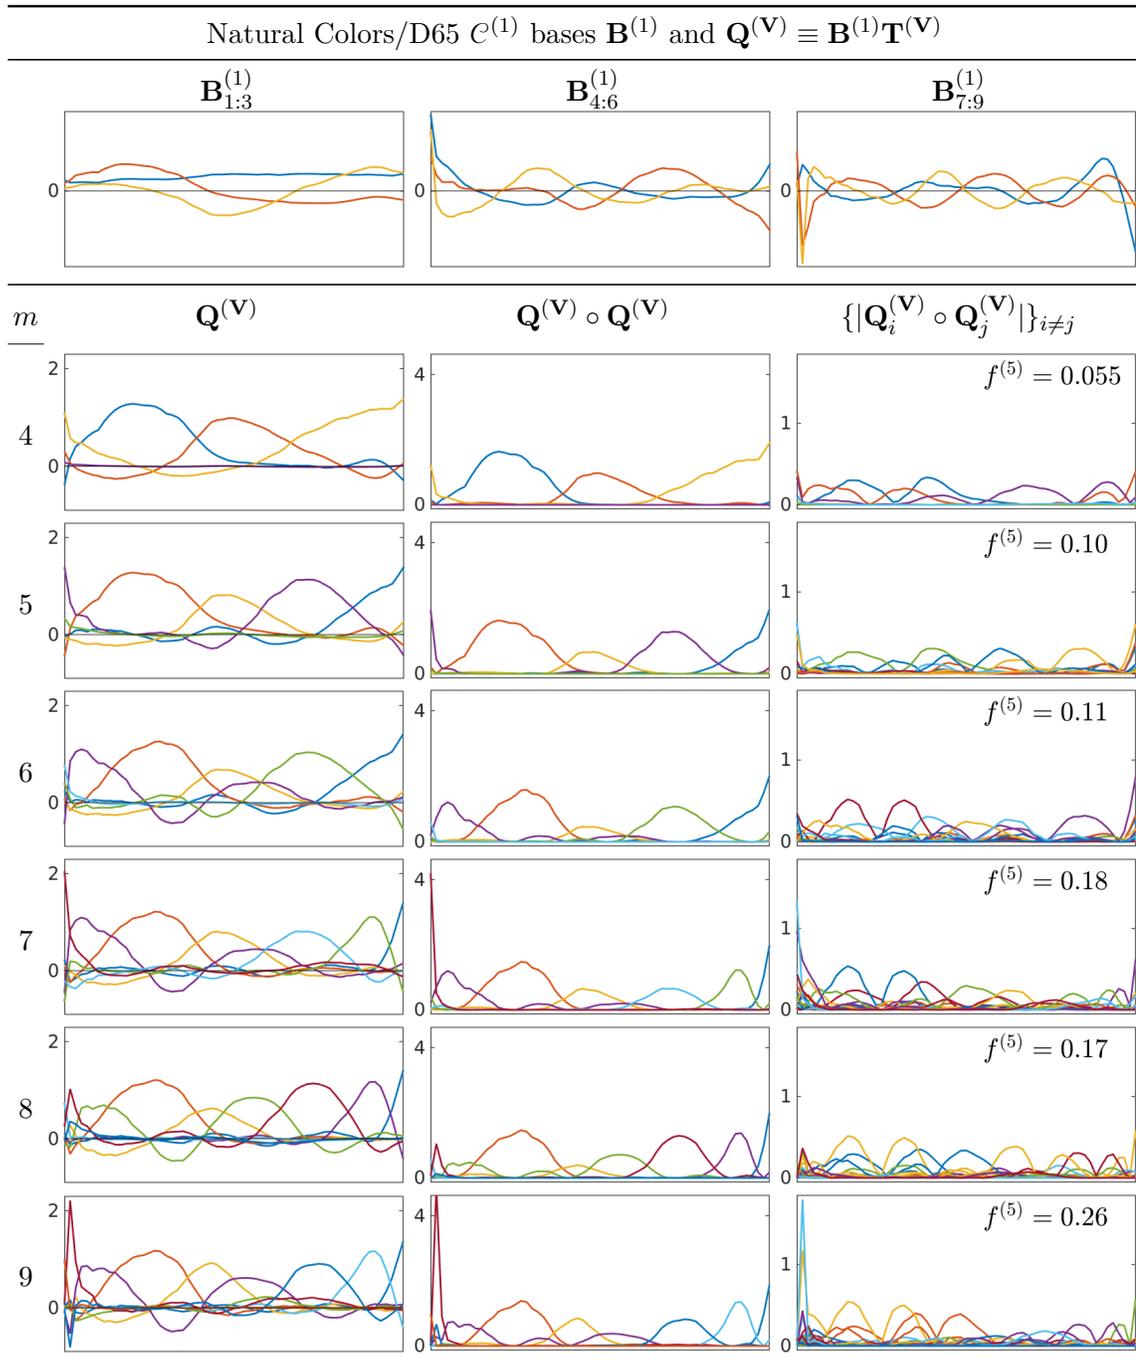


Figure 5.7: For Natural Colors $\mathcal{C}^{(1)}$ color signals formed with D65, the orthonormal basis $\mathbf{B}^{(1)}$, the sharp bases $\mathbf{Q}^{(\mathbf{V})} \equiv \mathbf{B}^{(1)}\mathbf{T}^{(\mathbf{V})}$ (for dimensionalities $m \equiv 4, \dots, 9$), their squares and the absolute values of the pairwise products of their columns.

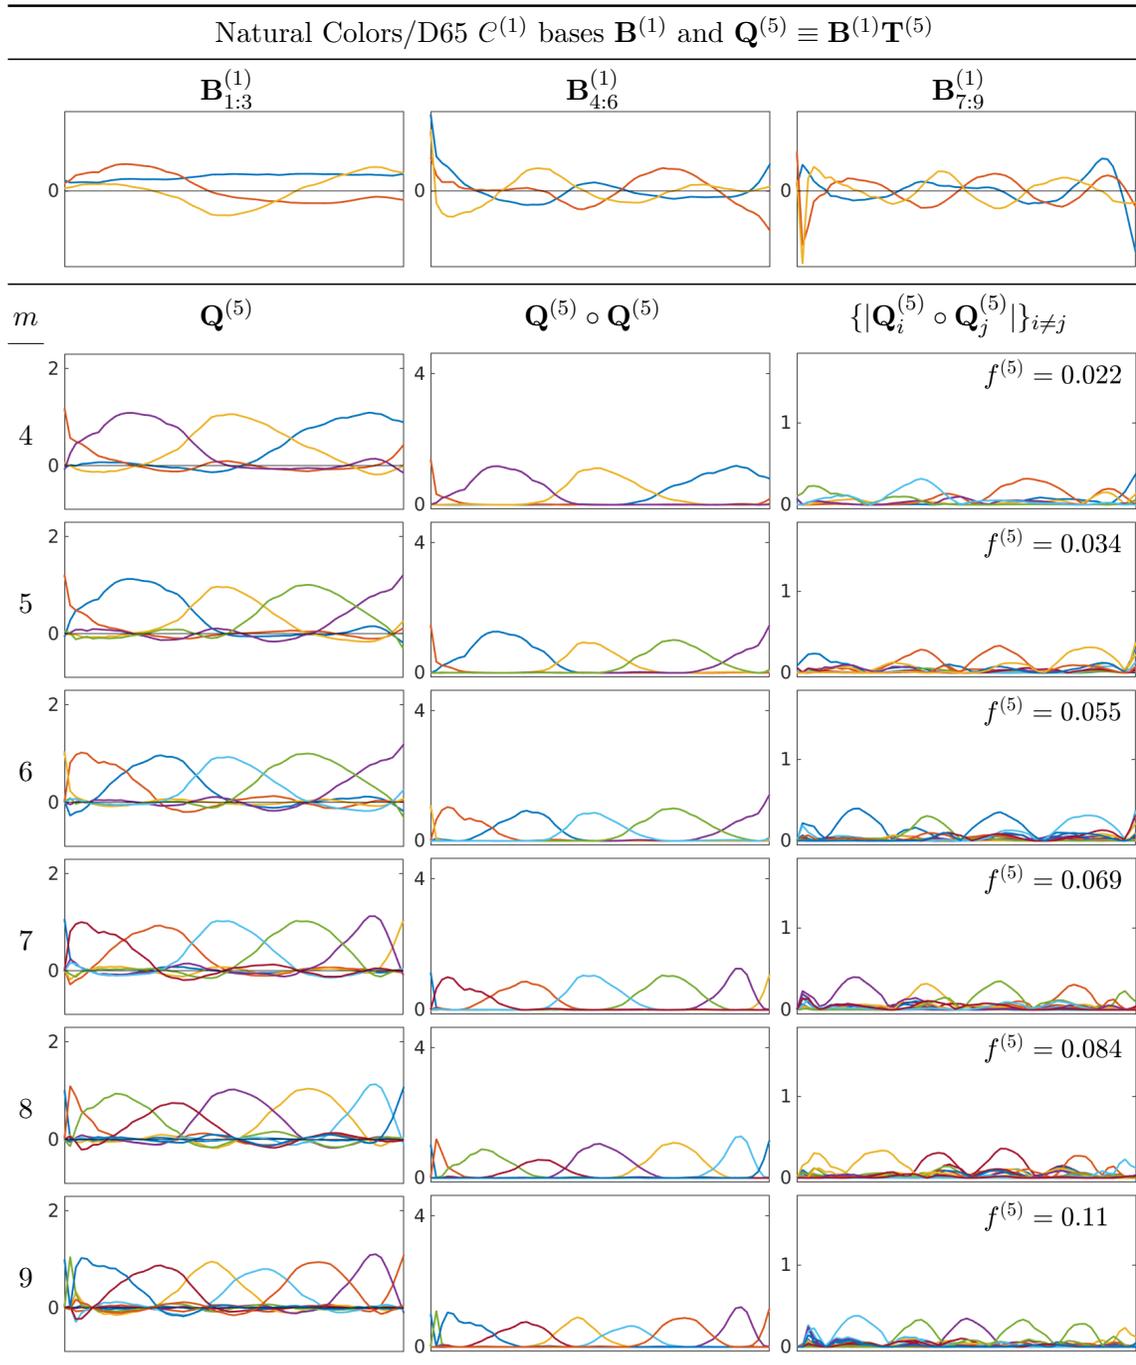


Figure 5.8: For Natural Colors $\mathcal{C}^{(1)}$ color signals formed with D65, the orthonormal basis $\mathbf{B}^{(1)}$, the sharp bases $\mathbf{Q}^{(5)} \equiv \mathbf{B}^{(1)}\mathbf{T}^{(5)}$ (for dimensionalities $m \equiv 4, \dots, 9$), their squares and the absolute values of the pairwise products of their columns.

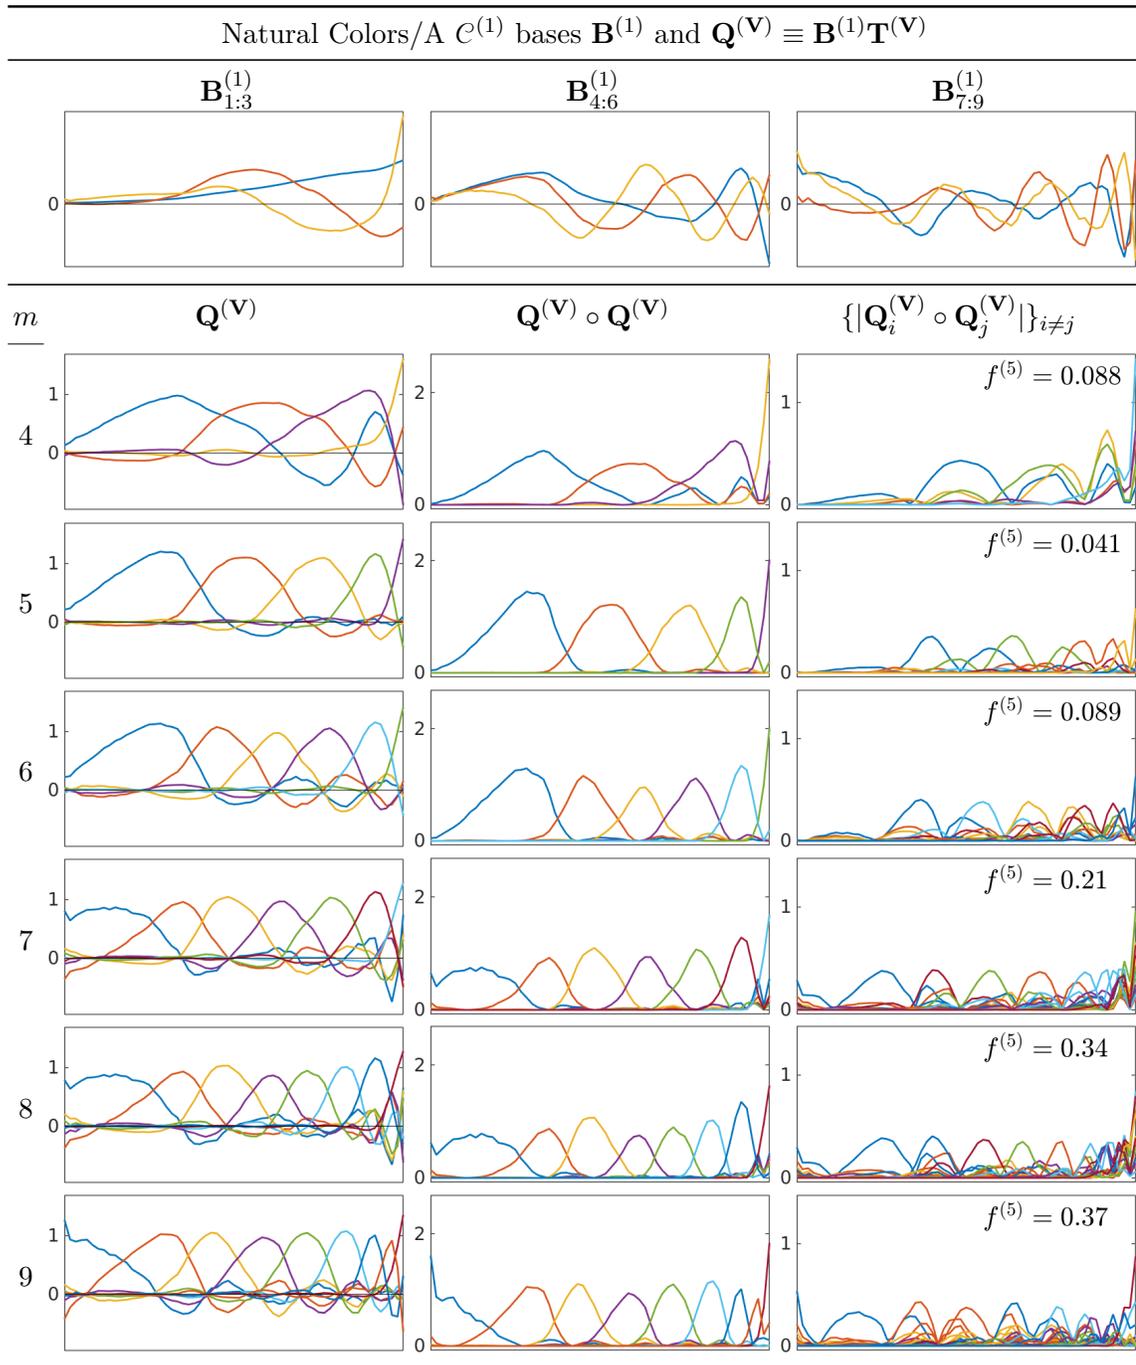


Figure 5.9: For Natural Colors $\mathcal{C}^{(1)}$ color signals formed with illuminant A, the orthonormal basis $\mathbf{B}^{(1)}$, the sharp bases $\mathbf{Q}^{(\mathbf{V})} \equiv \mathbf{B}^{(1)}\mathbf{T}^{(\mathbf{V})}$ (for dimensionalities $m \equiv 4, \dots, 9$), their squares and the absolute values of the pairwise products of their columns.

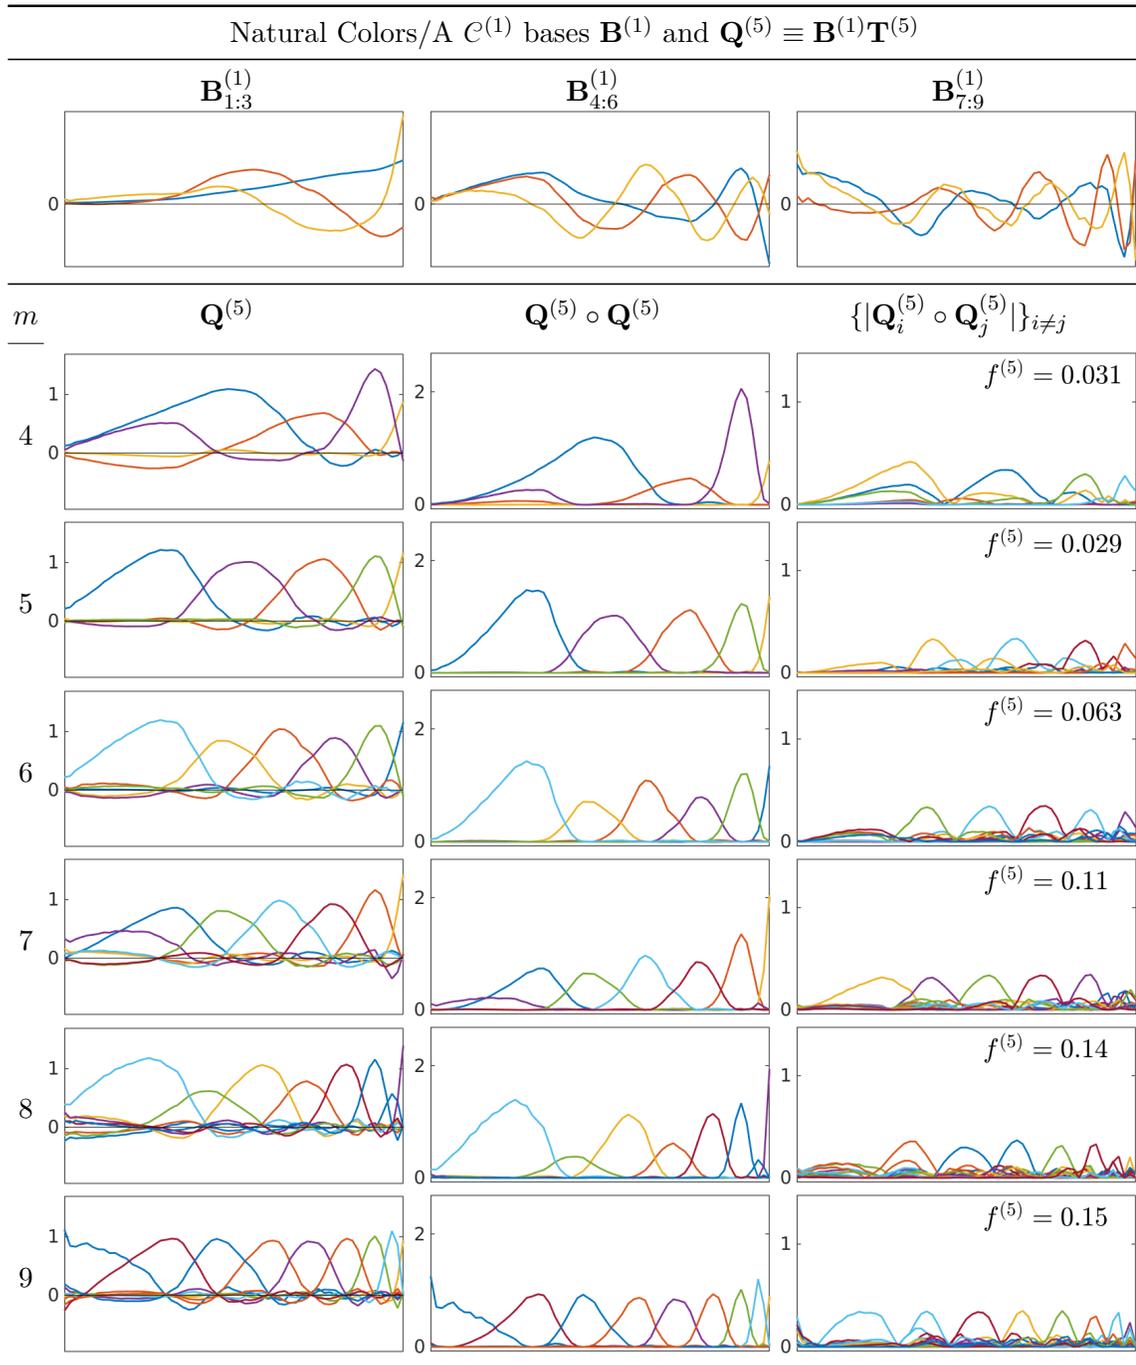


Figure 5.10: For Natural Colors $\mathcal{C}^{(1)}$ color signals formed with illuminant A, the orthonormal basis $\mathbf{B}^{(1)}$, the sharp bases $\mathbf{Q}^{(5)} \equiv \mathbf{B}^{(1)}\mathbf{T}^{(5)}$ (for dimensionalities $m \equiv 4, \dots, 9$), their squares and the absolute values of the pairwise products of their columns.

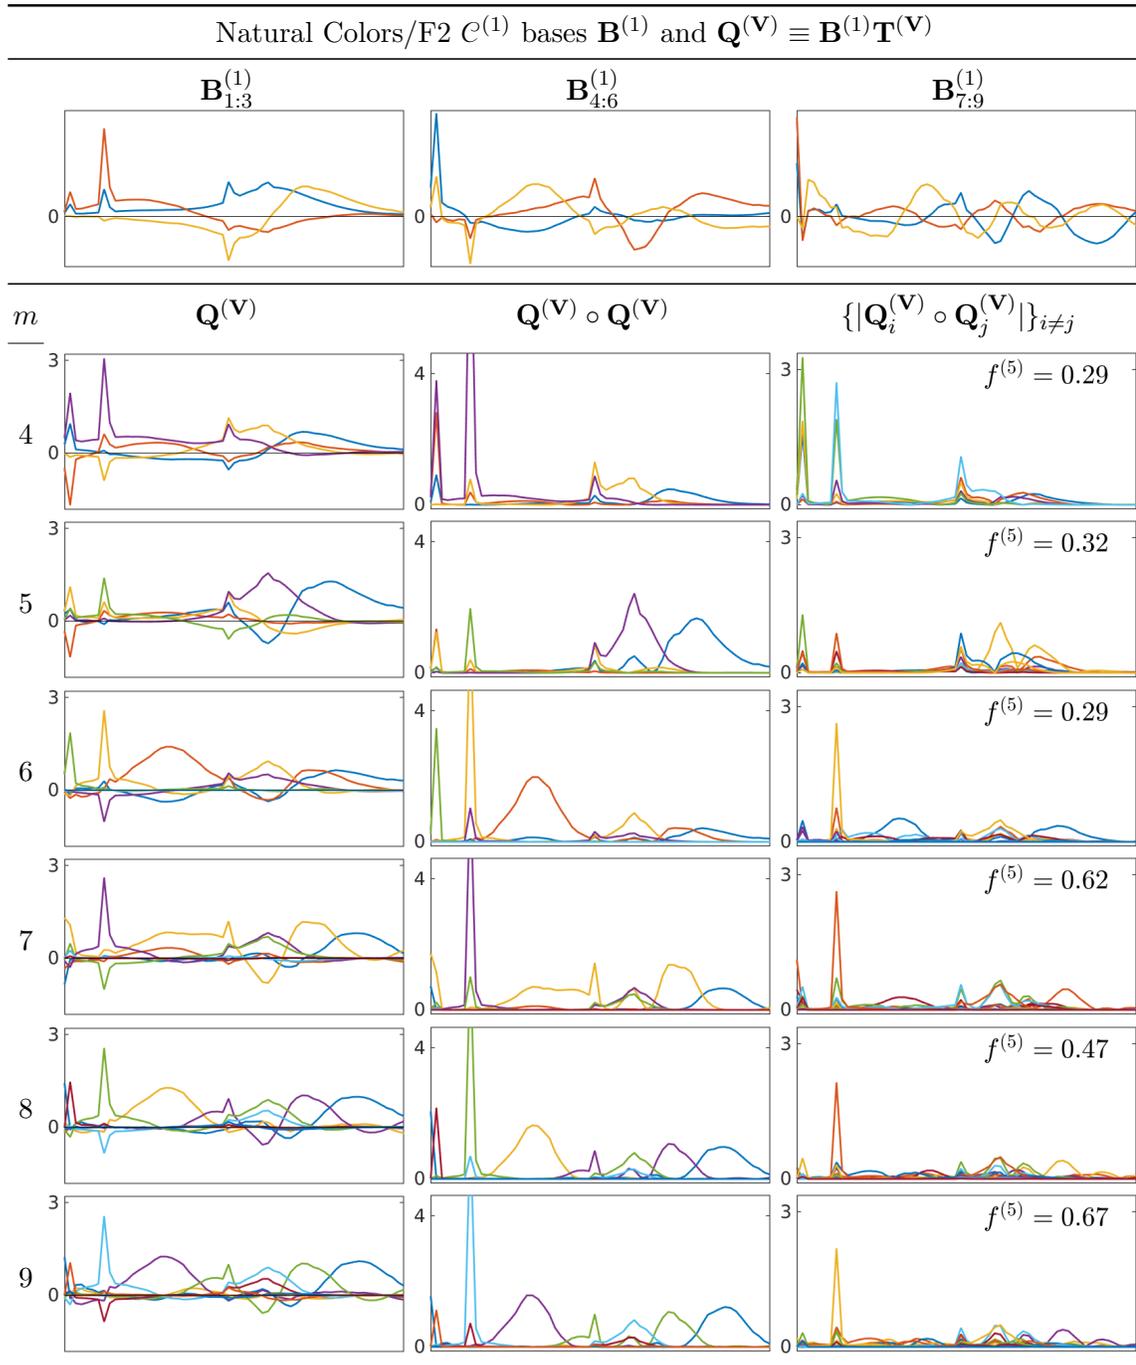


Figure 5.11: For Natural Colors $\mathcal{C}^{(1)}$ color signals formed with F2, the orthonormal basis $\mathbf{B}^{(1)}$, the sharp bases $\mathbf{Q}^{(\mathbf{V})} \equiv \mathbf{B}^{(1)}\mathbf{T}^{(\mathbf{V})}$ (for dimensionalities $m \equiv 4, \dots, 9$), their squares and the absolute values of the pairwise products of their columns.

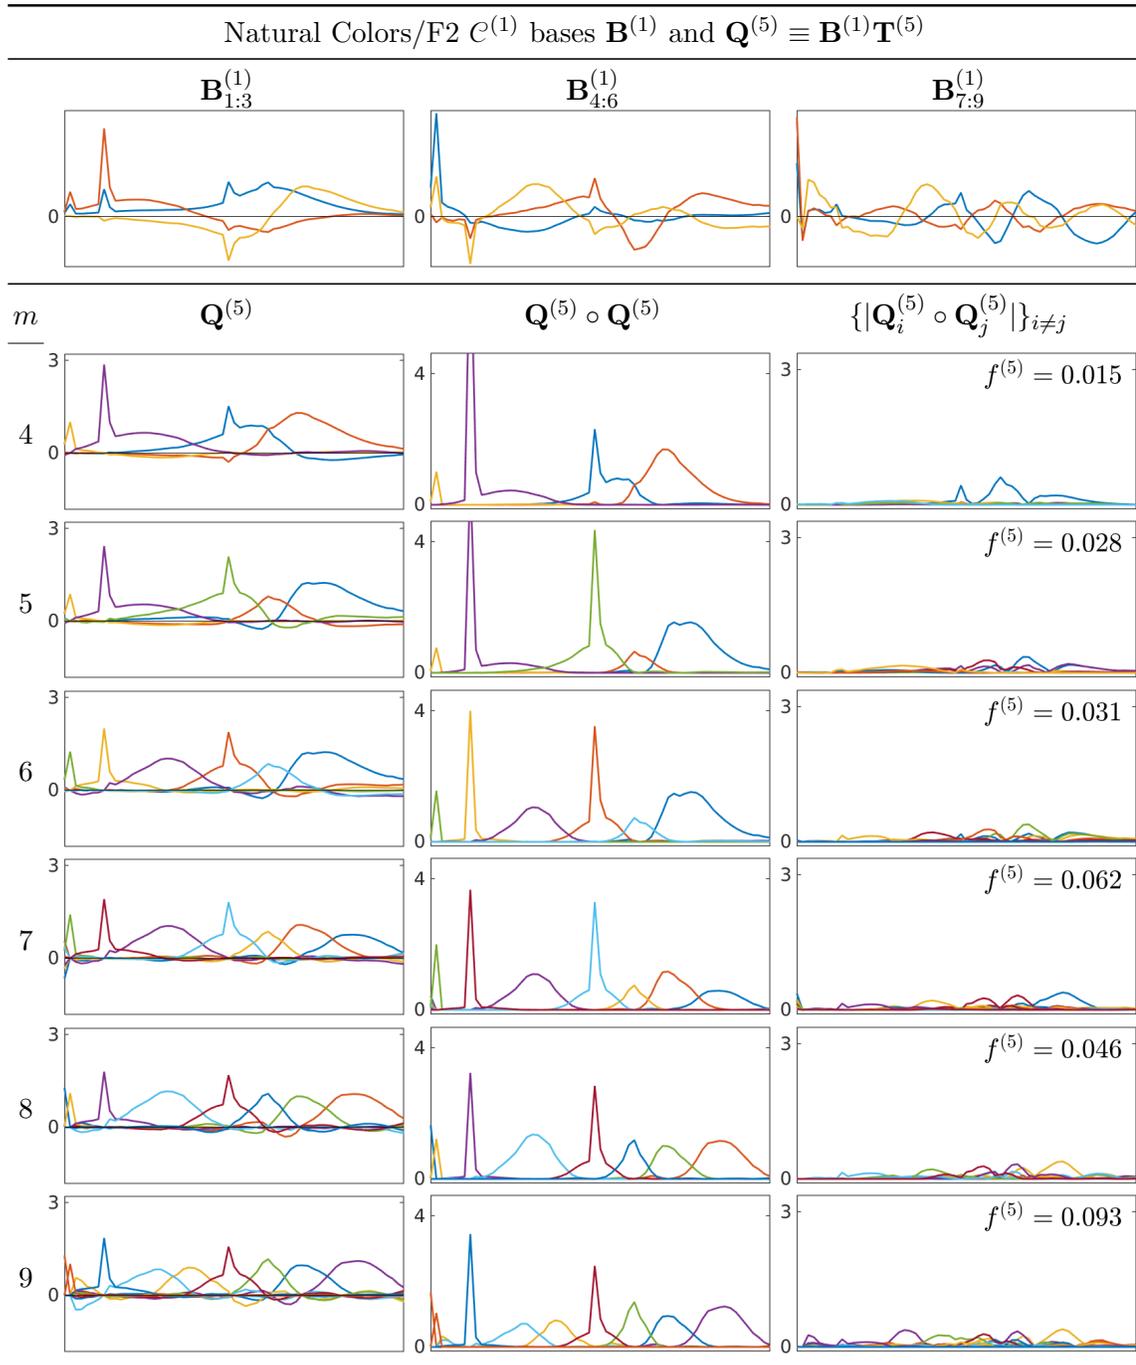


Figure 5.12: For Natural Colors $\mathcal{C}^{(1)}$ color signals formed with F2, the orthonormal basis $\mathbf{B}^{(1)}$, the sharp bases $\mathbf{Q}^{(5)} \equiv \mathbf{B}^{(1)}\mathbf{T}^{(5)}$ (for dimensionalities $m \equiv 4, \dots, 9$), their squares and the absolute values of the pairwise products of their columns.

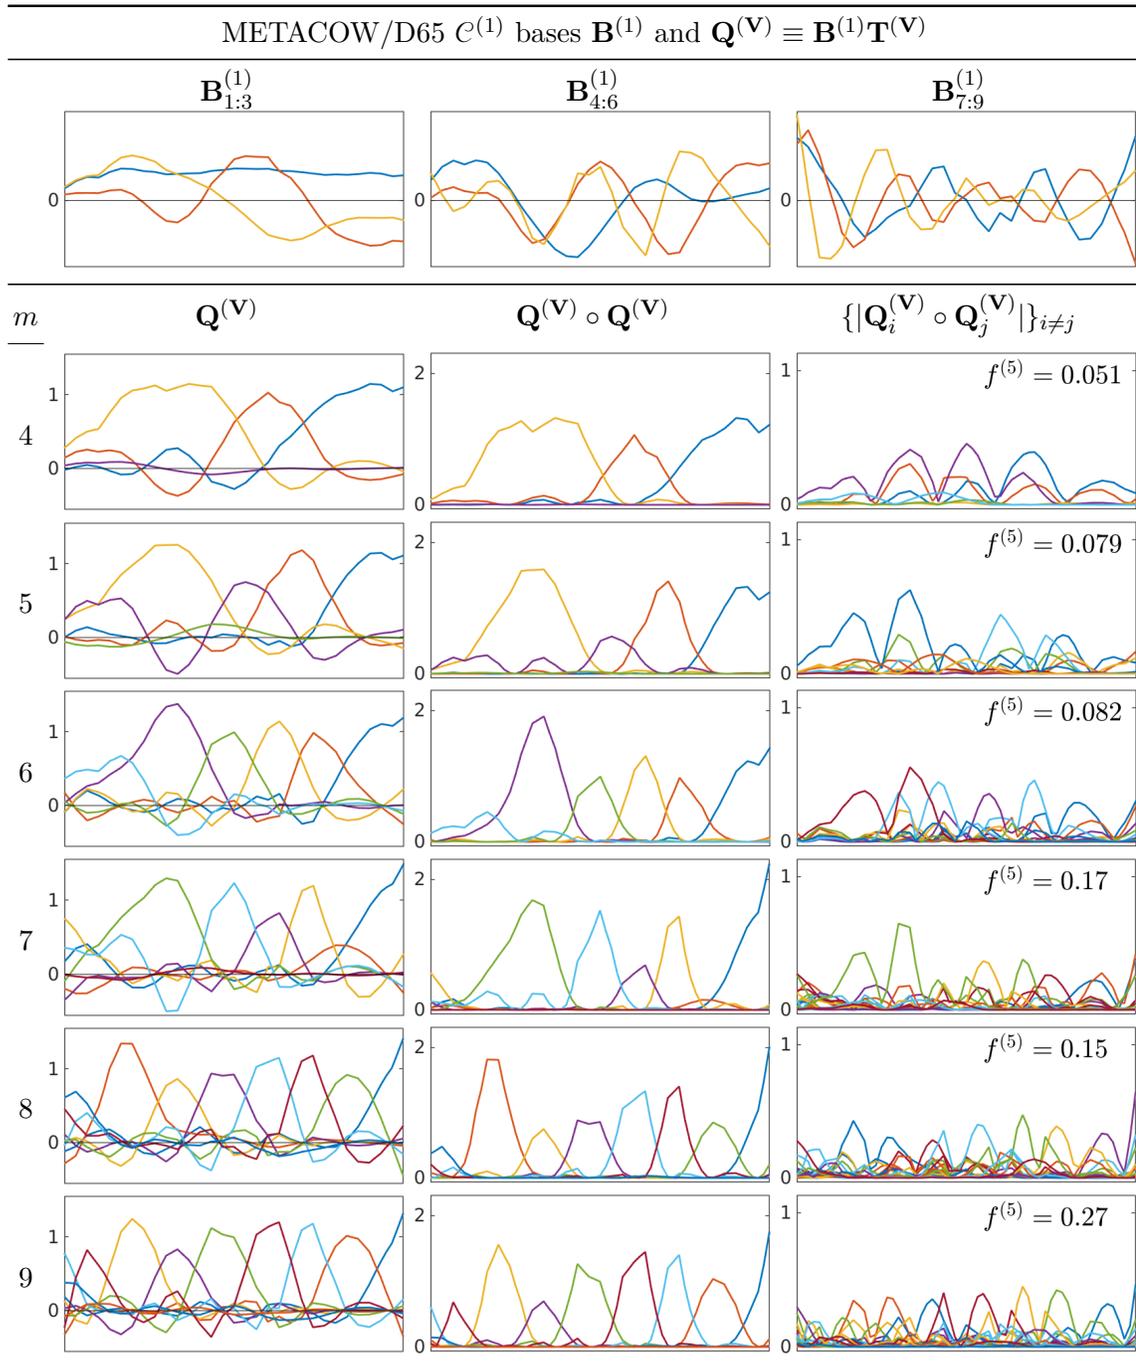


Figure 5.13: For METACOW $\mathcal{C}^{(1)}$ color signals formed with D65, the orthonormal basis $\mathbf{B}^{(1)}$, the sharp bases $\mathbf{Q}^{(\mathbf{V})} \equiv \mathbf{B}^{(1)}\mathbf{T}^{(\mathbf{V})}$ (for dimensionalities $m \equiv 4, \dots, 9$), their squares and the absolute values of the pairwise products of their columns.

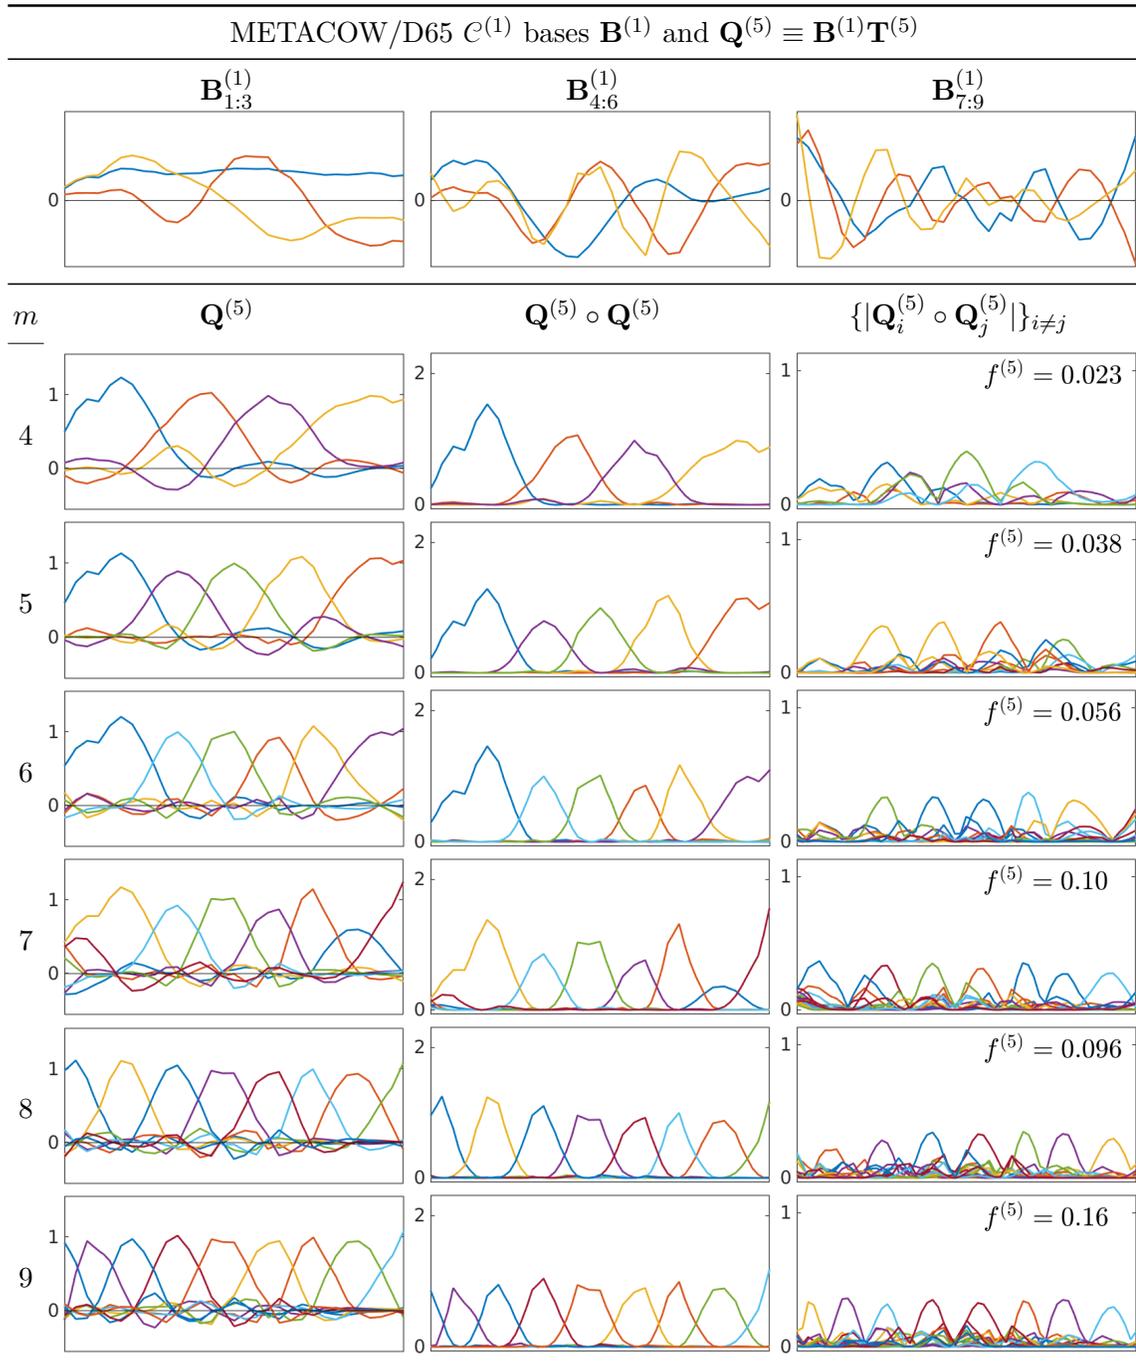


Figure 5.14: For METACOW $\mathcal{C}^{(1)}$ color signals formed with D65, the orthonormal basis $\mathbf{B}^{(1)}$, the sharp bases $\mathbf{Q}^{(5)} \equiv \mathbf{B}^{(1)}\mathbf{T}^{(5)}$ (for dimensionalities $m \equiv 4, \dots, 9$), their squares and the absolute values of the pairwise products of their columns.

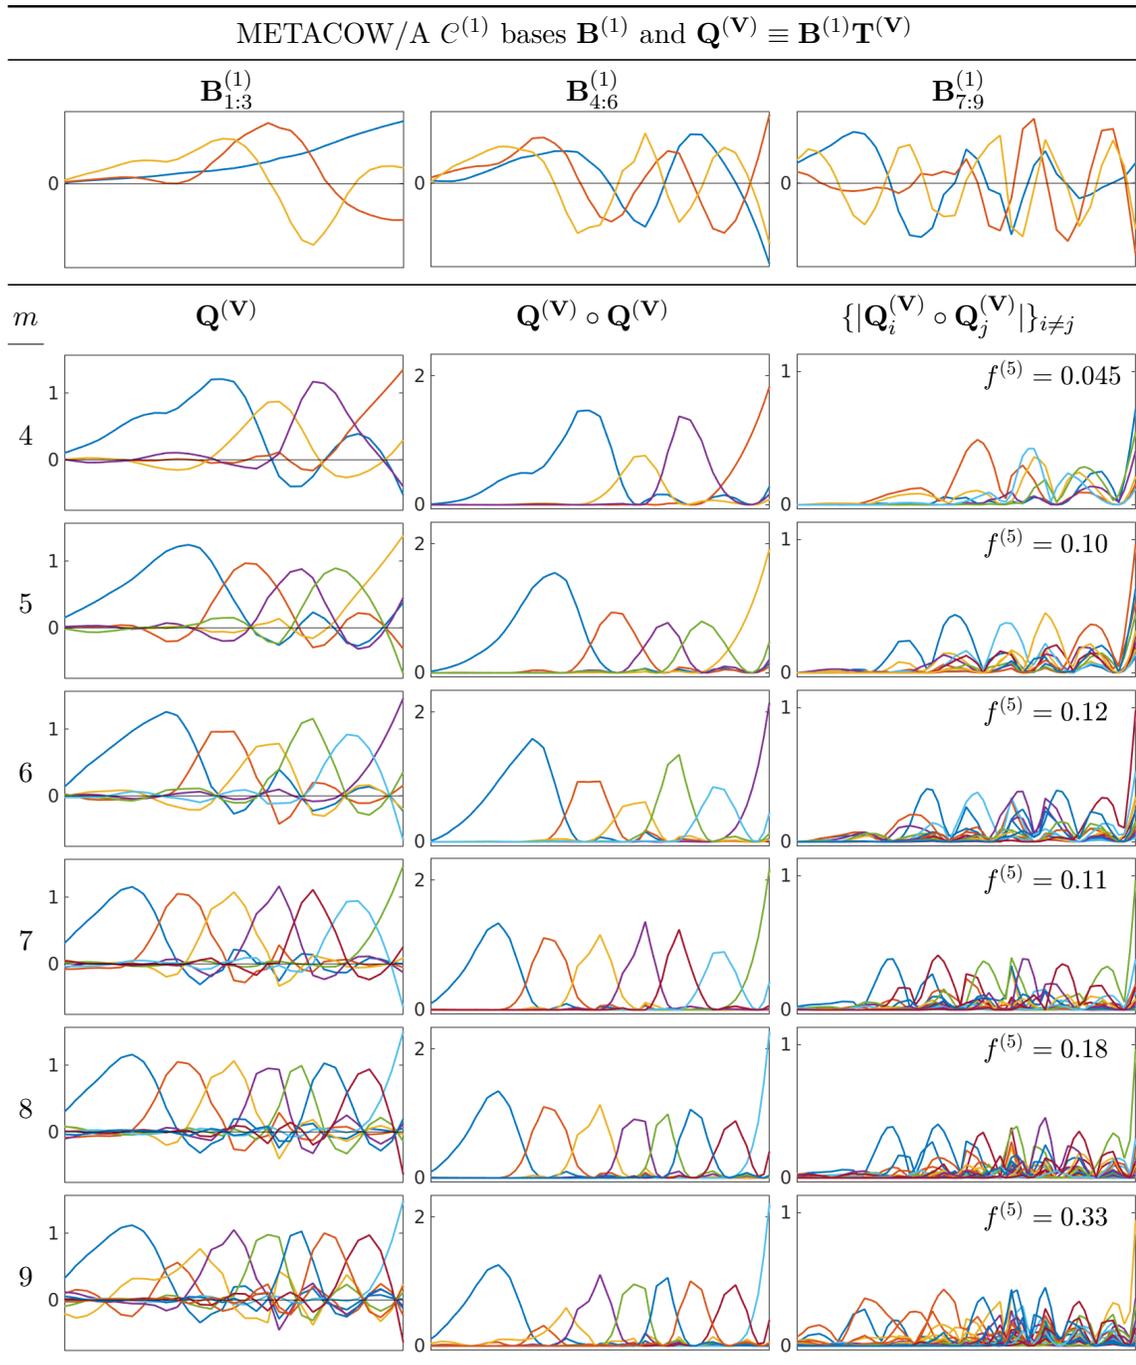


Figure 5.15: For METACOW $\mathcal{C}^{(1)}$ color signals formed with illuminant A, the orthonormal basis $\mathbf{B}^{(1)}$, the sharp bases $\mathbf{Q}^{(\mathbf{V})} \equiv \mathbf{B}^{(1)}\mathbf{T}^{(\mathbf{V})}$ (for dimensionalities $m \equiv 4, \dots, 9$), their squares and the absolute values of the pairwise products of their columns.

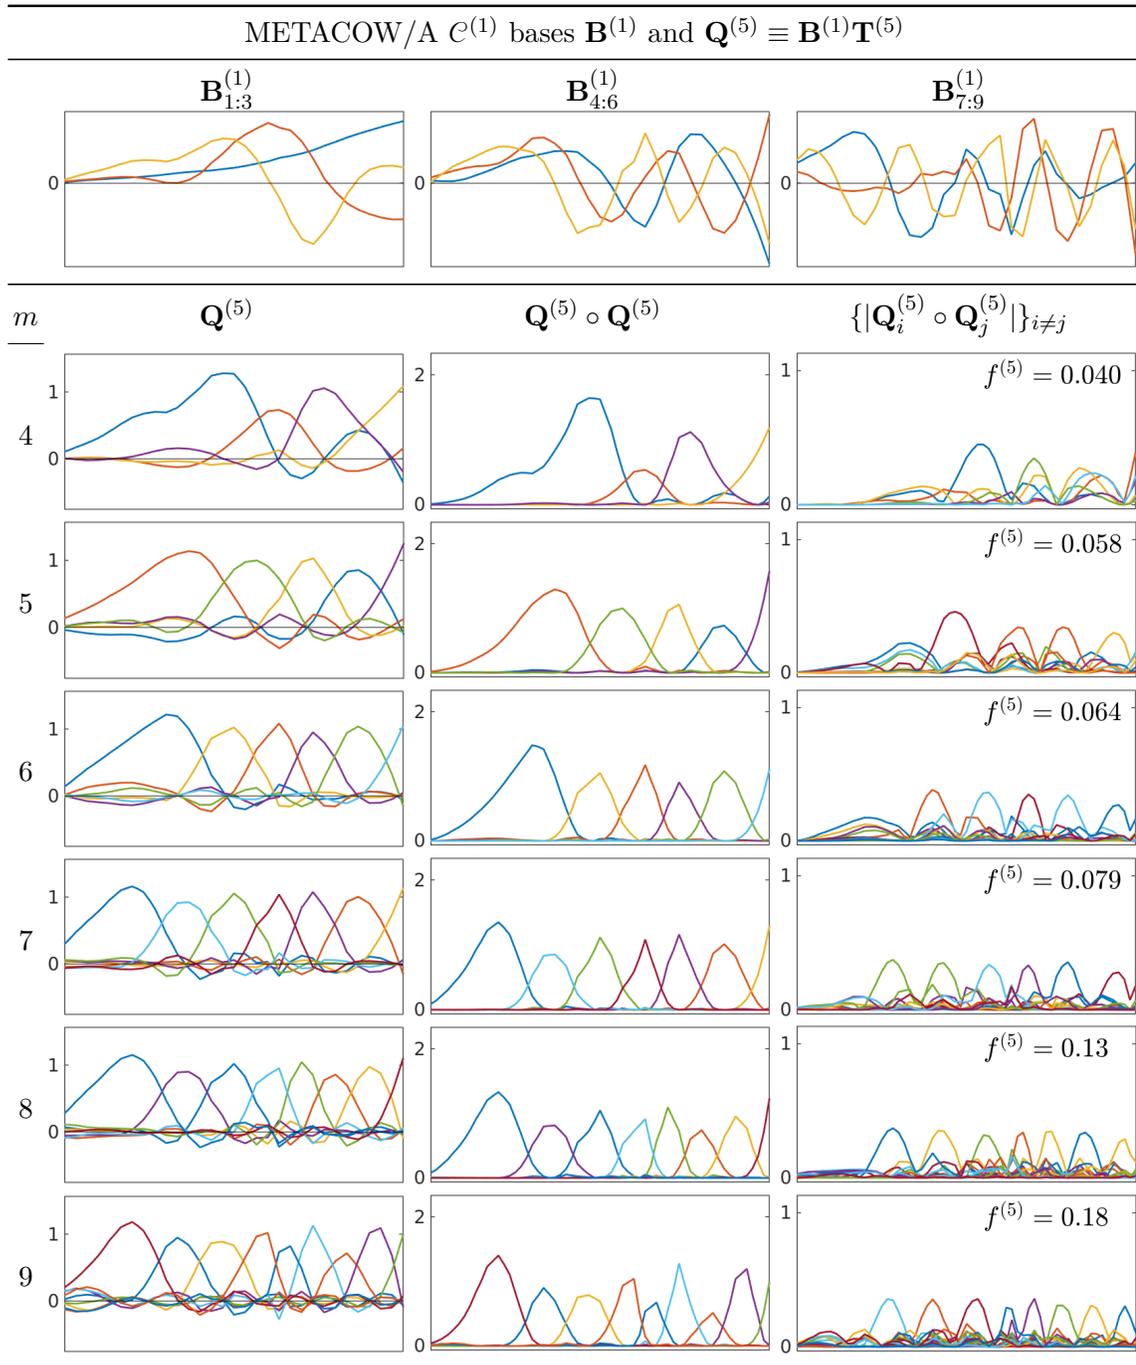


Figure 5.16: For METACOW $\mathcal{C}^{(1)}$ color signals formed with illuminant A, the orthonormal basis $\mathbf{B}^{(1)}$, the sharp bases $\mathbf{Q}^{(5)} \equiv \mathbf{B}^{(1)}\mathbf{T}^{(5)}$ (for dimensionalities $m \equiv 4, \dots, 9$), their squares and the absolute values of the pairwise products of their columns.

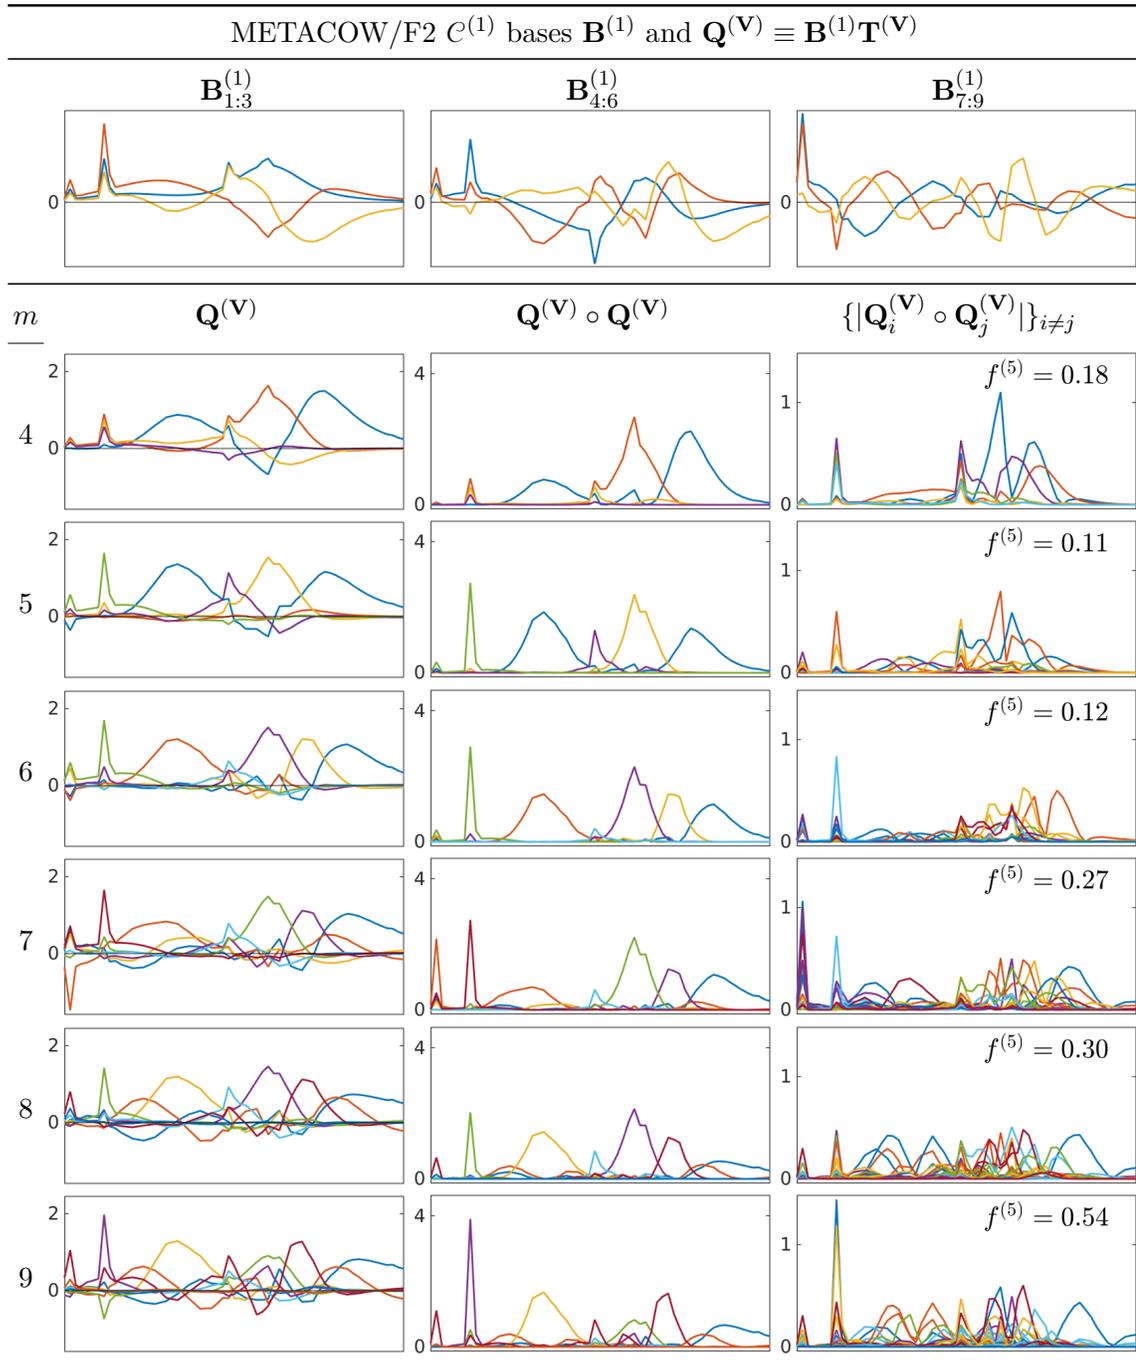


Figure 5.17: For METACOW $\mathcal{C}^{(1)}$ color signals formed with F2, the orthonormal basis $\mathbf{B}^{(1)}$, the sharp bases $\mathbf{Q}^{(\mathbf{V})} \equiv \mathbf{B}^{(1)}\mathbf{T}^{(\mathbf{V})}$ (for dimensionalities $m \equiv 4, \dots, 9$), their squares and the absolute values of the pairwise products of their columns.

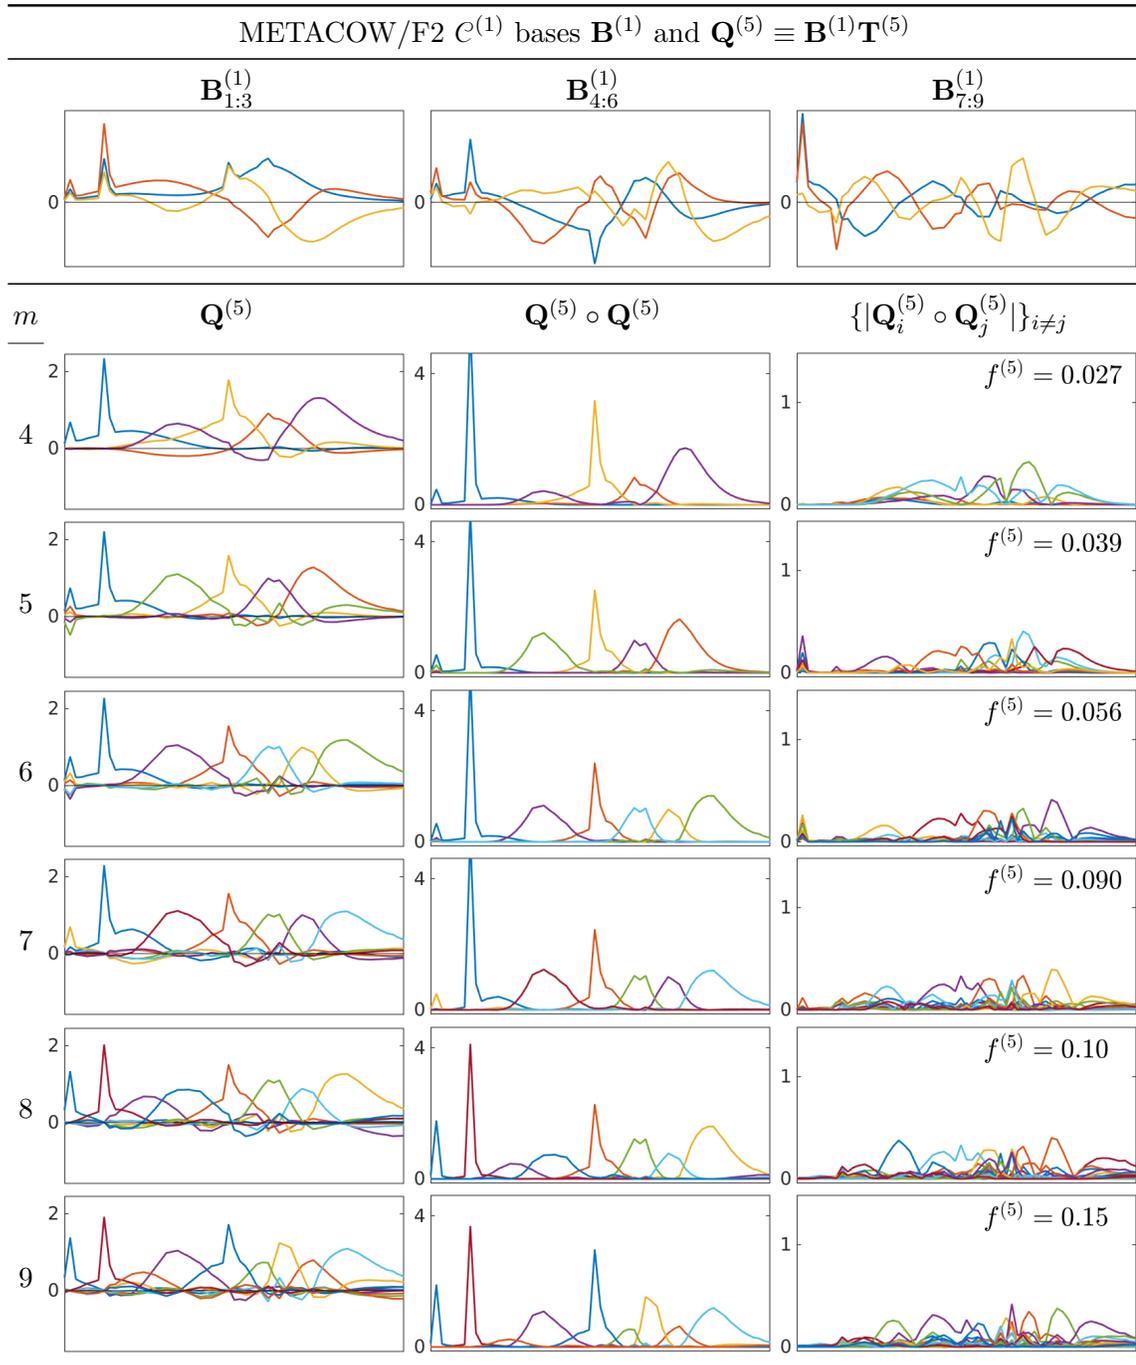


Figure 5.18: For METACOW $\mathcal{C}^{(1)}$ color signals formed with F2, the orthonormal basis $\mathbf{B}^{(1)}$, the sharp bases $\mathbf{Q}^{(5)} \equiv \mathbf{B}^{(1)}\mathbf{T}^{(5)}$ (for dimensionalities $m \equiv 4, \dots, 9$), their squares and the absolute values of the pairwise products of their columns.

approximation error, or by solving a simpler subproblem yielding an approximate solution by implicitly minimizing an upper bound on it. The simplest of these, usually yielding the largest approximation error, is spectral sharpening. All of these problems and methods for solving them are described in the next four sections, starting with the largest to smallest subproblems before discussing $f^{(1)}$.

5.2 By minimizing the residual $\|\Delta\|_F^2$

Because $\|\Delta\mathbf{c}\| \leq \|\Delta\|\|\mathbf{c}\|$, where $\mathbf{c} \equiv \mathbf{c}^{(1)}$, and because \mathbf{c} does not depend on \mathbf{T} , an approximate solution to (5.1) is obtained by solving

$$\mathbf{T}^{(2)} \equiv \arg \min_{\mathbf{T}} f^{(2)}(\mathbf{T}, \mathbf{B}, \widehat{\mathbf{U}}^{(s)}), \text{ where} \quad (5.11)$$

$$f^{(2)} \equiv \|\Delta\|_F^2 \equiv \sum_{i=1}^n \|\mathbf{B}^T \text{diag}(\widehat{\mathbf{U}}_i^{(s)})\mathbf{B} - \mathbf{T} \text{diag}(\widetilde{\mathbf{u}}_i^{(s)})\mathbf{T}^{-1}\|_F^2. \quad (5.12)$$

Like (5.1), this is an unconstrained, nonlinear least squares problem with m^2 parameters and $n \cdot m^2$ squared residuals. Many iterative methods exist for solving this type of problem [61, 133], including quasi-Newton methods, such as the BFGS algorithm [15, 60, 68, 154], or methods designed specifically for nonlinear least-squares problems, such as the Levenberg-Marquardt algorithm [108, 120].

5.3 By approximate joint diagonalization of $\{\mathbf{R}^{(i)}\}_{i=1}^n$

Optimizations similar to (5.11) appear in some signal processing problems, such as blind source separation (BSS) used to recover a set of unknown source signals from an observed mixture [37]. One approach to BSS involves the approximate joint diagonalization (AJD) of a set of square target matrices expressing a statistical property of the mixture, such as auto-correlation [24]. AJD algorithms seek a matrix \mathbf{T} (often assumed orthogonal, so that $\mathbf{T}^{-1} = \mathbf{T}^T$) that by some criterion can best approximately diagonalize the target matrices, which, unless they all commute, cannot be diagonalized exactly when there are more than two. With \mathbf{T} constrained to avoid the trivial solution, a common criterion is the sum of the squared off-diagonal elements of the transformed target matrices. Posed as an AJD problem, an approximate solution

to (5.11) is thus

$$\mathbf{T}^{(3)} \equiv \mathbf{T} \operatorname{diag}(\mathbf{T}^{-1} \mathbf{B}^T \mathbf{1}), \text{ where} \quad (5.13)$$

$$\mathbf{T} \equiv \arg \min_{\substack{\mathbf{T} \\ \text{s.t. } \mathbf{T}^T \mathbf{T} = \mathbf{I} \text{ or} \\ \operatorname{diag}(\mathbf{T}^T \mathbf{T}) = \mathbf{1}}} f^{(3)}(\mathbf{T}, \mathbf{B}, \widehat{\mathbf{U}}^{(\mathcal{S})}), \text{ and} \quad (5.14)$$

$$\begin{aligned} f^{(3)} &\equiv \sum_{i=1}^n \|\operatorname{off}[\mathbf{T}^{-1} \mathbf{B}^T \operatorname{diag}(\widehat{\mathbf{U}}_i^{(\mathcal{S})}) \mathbf{B} \mathbf{T}]\|_F^2 = \sum_{i=1}^n \|\operatorname{off}(\mathbf{Q}^+ \operatorname{diag}(\widehat{\mathbf{U}}_i^{(\mathcal{S})}) \mathbf{Q})\|_F^2 \\ &\equiv \sum_{i=1}^n \|\mathbf{Q}^+ \operatorname{diag}(\widehat{\mathbf{U}}_i^{(\mathcal{S})}) \mathbf{Q} - \operatorname{ddiag}[\mathbf{Q}^+ \operatorname{diag}(\widehat{\mathbf{U}}_i^{(\mathcal{S})}) \mathbf{Q}]\|_F^2. \end{aligned}$$

Generic algorithms can find solutions to (5.14), a nonlinear least squares problem with quadratic constraints, but methods that exploit its special structure will find them more efficiently. These methods can be categorized by the constraints placed on solutions. In general, constraining \mathbf{T} to be orthogonal results in faster convergence, but looser restrictions can yield lower objective function values. An example of the former is a Jacobi-like algorithm described by Cardoso and Souloumiac that constructs \mathbf{T} as the product of a sequence of Givens rotations that reduce the sum of the squared target matrices' off-diagonal elements on each iteration [21]. An efficient example of the latter is the Uniformly Weighted Approximate Joint Diagonalization (UWAJD) algorithm described by Tichavský et al. [161].

A difference between solutions to (5.11) and (5.14), and a source of additional approximation error, is that the diagonal elements of $\mathbf{Q}^+ \operatorname{diag}(\widehat{\mathbf{U}}_i^{(\mathcal{S})}) \mathbf{Q}$ are ignored in (5.14), while the corresponding $\widetilde{\mathbf{u}}_i^{(\mathcal{S})} \equiv \mathbf{Q}^+ \widehat{\mathbf{U}}_i^{(\mathcal{S})}$ in (5.11) are not. That these, for sufficiently sharp \mathbf{Q} , are approximately equal can be seen by noting (5.13), a normalization that yields $\mathbf{Q}^+ \mathbf{1} = [\mathbf{Q}^{+T} \circ \mathbf{Q}]^T \mathbf{1} = \mathbf{1}$, and recalling that \mathbf{Q}^{+T} and \mathbf{Q} tend to be sharpened simultaneously within approximately the same wavelength intervals. (In the limit, where disjoint $\mathbf{Q}_i \circ (\mathbf{Q}^{+T})_j = \mathbf{0}$, each \mathbf{Q}_i and $(\mathbf{Q}^{+T})_j$ assumes the form of a box function with \mathbf{Q}_i of unity height, and $\mathbf{Q} \circ \mathbf{Q}^{+T} = \mathbf{Q}^{+T}$.) Thus, \mathbf{Q}^{+T} approximates $\mathbf{Q}^{+T} \circ \mathbf{Q}$ (see Fig. 5.19 for two examples), and

$$[\mathbf{Q}^+ \operatorname{diag}(\widehat{\mathbf{U}}_i^{(\mathcal{S})}) \mathbf{Q}]_{jj} = \langle \mathbf{Q}_j^{+T} \circ \mathbf{Q}_j, \widehat{\mathbf{U}}_i^{(\mathcal{S})} \rangle \approx \langle \mathbf{Q}_j^{+T}, \widehat{\mathbf{U}}_i^{(\mathcal{S})} \rangle \equiv \widetilde{\mathbf{u}}_{ji}^{(\mathcal{S})}.$$

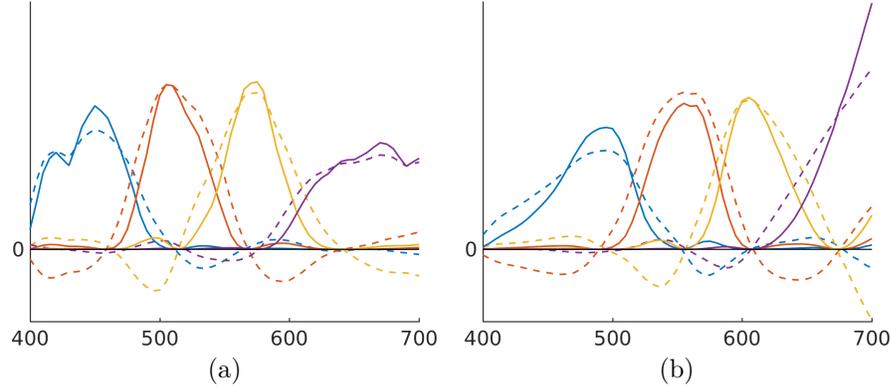


Figure 5.19: $\mathbf{Q}^{+T} \circ \mathbf{Q}$ (solid lines) and \mathbf{Q}^{+T} (dashed lines) for Munsell $\mathcal{C}^{(2)}$ SPDs formed with illuminants (a) D65 and (b) A, where $\mathbf{Q} \equiv \mathbf{B}^{(2)}\mathbf{T}$, and \mathbf{T} approximately diagonalizes $\{\mathbf{R}^{(i)}\}_{i=1}^n$.

5.4 By sharpening \mathbf{Q}

In [49] Drew and Finlayson find \mathbf{T}^{-T} by solving m optimization problems, each of which requires a user to choose a subset of wavelengths $\mathcal{L}^{(i)} \subset \{400, 405, \dots, 700\}$, or *sharpening interval*, in which the corresponding i -th column of \mathbf{Q}^{+T} is sharpened:

$$\mathbf{T}^{(4)} \equiv \mathbf{T}^{-T} \text{diag}(\mathbf{T}^T \mathbf{B}^T \mathbf{1}), \text{ where} \quad (5.15)$$

$$\mathbf{T} \equiv \underset{\mathbf{T}}{\text{arg min}} \quad f^{(4)}(\mathbf{T}, \mathbf{B}, \mathcal{L}) \quad (5.16)$$

s.t. $\text{diag}(\mathbf{T}^T \mathbf{T})^{-1} = \mathbf{1}$

$$f^{(4)} \equiv \sum_{i=1}^m \|\mathbf{B}(\mathcal{L}^{(i)})(\mathbf{T}^{-T})_i\|_2^2,$$

and $\mathbf{B}(\mathcal{L}^{(i)})$ is the basis \mathbf{B} with entries in rows corresponding to $\lambda \notin \mathcal{L}^{(i)}$ set to zero. Claiming that it “makes intuitive sense,” the authors constrain the elements of \mathbf{Q}^+ to be nonnegative, but this is mathematically unnecessary. In an earlier article on spectral sharpening [48], the same authors show that without this constraint (5.16) is equivalent to m eigenvector problems, each of which sets the i -th column of \mathbf{T}^{-T} to the eigenvector $\mathbf{t}^{(k)}$ of $\mathbf{B}(\mathcal{L}^{(i)})^T \mathbf{B}(\mathcal{L}^{(i)})$ that minimizes $\|\mathbf{B}(\overline{\mathcal{L}^{(i)}})\mathbf{t}^{(k)}\|_2$, where $\overline{\mathcal{L}^{(i)}} \equiv \{400, 405, \dots, 700\} \setminus \mathcal{L}^{(i)}$ is the relative complement of $\mathcal{L}^{(i)}$. Because better results are obtained without it, the nonnegativity constraint on \mathbf{Q}^+ is omitted when using this method to obtain the results reported in the next chapter.

Whether the sharpest basis yields the lowest approximation error is a question left unaddressed by Drew and Finlayson. Likewise, which set of sharpening intervals

$\mathcal{L} \equiv \{\mathcal{L}^{(i)}\}_{i=1}^m$ yields the sharpest basis is itself an optimization problem that the authors neither pose nor solve, suggesting simply, without justification, that they be connected (i.e., that intervals do not contain gaps) and that they have approximately uniform size and even spacing between 400 and 700 nm. If gaps are permitted between intervals, as in [49], and each interval is allowed to be any size but required to be connected, there are $\binom{n}{2m}$ possible choices, which makes exhaustive searches impractical. Using metaheuristic methods [159] to sample the parameter space stochastically (for example, by simulated annealing [97] or by particle swarm optimization [51]), it can be shown that intervals yielding the sharpest bases tend to be neither uniform nor evenly spaced nor even necessarily connected [170]. The same is also true of intervals that yield the lowest approximation error, which in general do not yield the sharpest bases. In other words, the bases found by the method of Drew and Finlayson are neither the sharpest nor optimal.

A choice of wavelengths can be avoided altogether by posing a single optimization problem that maximizes the sharpness of the *whole* basis by minimizing the overlap of the transformed basis vectors across all wavelengths:

$$\mathbf{T}^{(5)} \equiv \mathbf{T} \operatorname{diag}(\mathbf{T}^{-1} \mathbf{B}^T \mathbf{1}), \text{ where} \quad (5.17)$$

$$\mathbf{T} \equiv \underset{\mathbf{T}}{\operatorname{arg\,min}} \quad f^{(5)}(\mathbf{T}, \mathbf{B}) \text{ and} \quad (5.18)$$

s.t. $\operatorname{diag}(\mathbf{T}^T \mathbf{T}) = \mathbf{1}$

$$f^{(5)} \equiv \sum_{i \neq j} \langle \mathbf{Q}_i \circ \mathbf{Q}_i, \mathbf{Q}_j \circ \mathbf{Q}_j \rangle = \sum_{i \neq j} \|\mathbf{Q}_i \circ \mathbf{Q}_j\|_F^2.$$

This quadratically constrained, nonlinear objective function, which sums the scalar products of all pairs of squared columns of \mathbf{Q} , is similar to a family of problems in factor analysis, known as *rotation* problems, which seek to maximize the variance of a set of transformed vectors [126]. In particular, (5.18) is identical to the *quartimin*⁴ problem [85] (so named because it minimizes a quartic function), except for its constraint $\operatorname{diag}(\mathbf{T}^{-1} \mathbf{T}^{-T}) = \mathbf{1}$, which in our formulation would, before its normalization, constrain \mathbf{Q}^+ instead of \mathbf{Q} . Jennrich and Sampson describe an efficient method for solving this problem, which cycles through pairs of vectors, “rotating” one within their plane by a linear combination that minimizes $f^{(5)}$, until convergence.

⁴More precisely, the problem described by Jennrich and Sampson in [85] is named the *direct quartimin* problem to distinguish it from the original quartimin problem described by Carroll in [23], which seeks to transform a so-called *reference structure* related to the target vectors, rather than the target vectors themselves.

Factor analysis methods are categorized by the constraint placed on the *rotation matrix* \mathbf{T} . As in AJD, orthogonal constraints generally results in more efficient algorithms, but non-orthogonal transformations⁵, such as those sought by the quartimin problem, may be preferred, as they may yield better objective function values. An example of the former is the simpler *quartimax* problem (so named because it *maximizes* a quartic function), which is derived from the observation that because the row sums of $\mathbf{Q} \circ \mathbf{Q}$ are constant under orthogonal transformation the same is true of the sum of the row sums' squares. Thus, for $\mathbf{T}^T \mathbf{T} = \mathbf{I}$, given that

$$\sum_i \left(\sum_j \mathbf{Q}_{ij}^2 \right)^2 = \sum_{ij} \mathbf{Q}_{ij}^4 + \sum_{i,j \neq k} \mathbf{Q}_{ij}^2 \mathbf{Q}_{ik}^2 = \|\mathbf{Q} \circ \mathbf{Q}\|_F^2 + \sum_{i \neq j} \|\mathbf{Q}_i \circ \mathbf{Q}_j\|_F^2 = \text{constant},$$

a problem equivalent to the quartimin problem with orthogonal constraints, involving only m instead of $m(m-1)/2$ transformed vectors, is

$$\mathbf{T}^{(6)} \equiv \mathbf{T} \text{diag}(\mathbf{T}^{-1} \mathbf{B}^T \mathbf{1}), \text{ where} \quad (5.19)$$

$$\mathbf{T} \equiv \arg \max_{\substack{\mathbf{T} \\ \text{s.t. } \mathbf{T}^T \mathbf{T} = \mathbf{I}}} f^{(6)}(\mathbf{T}, \mathbf{B}) \equiv \sum_{ij} \mathbf{Q}_{ij}^4 = \|\mathbf{Q} \circ \mathbf{Q}\|_F^2. \quad (5.20)$$

This well studied problem, first described in 1954 by Neuhaus and Wrigley [132], has a long history in factor analysis. It was later shown by Kaiser [91] to be equivalent to three other quartic objective functions proposed at around the same time [23, 153, 59]. Neuhaus and Wrigley describe an efficient algorithm for solving it, which, much like the algorithm of Jennrich and Sampson for finding quartimin solutions, sequentially rotates pairs of vectors to maximize $f^{(6)}$ until convergence. By contrast, matrix methods optimize a measure of the whole transformed matrix \mathbf{Q} in each iteration [76, pp. 428–429], [155]. A simple algorithm described by Jennrich [84] converges to a solution by iteratively projecting the current solution in the direction of its gradient onto the manifold of orthogonal matrices. In each iteration, for the current $\mathbf{Q}^{(k)} \equiv \mathbf{B} \mathbf{T}^{(k)}$, the matrix $\mathbf{T}^{(k+1)}$ is sought that maximizes

$$\tilde{f}(\mathbf{T}^{(k+1)}) \equiv f(\mathbf{T}^{(k)}) + \text{tr} [\nabla f(\mathbf{T}^{(k)T})(\mathbf{T}^{(k+1)} - \mathbf{T}^{(k)})],$$

⁵Although *rotation* is mathematically defined as an orthogonal transformation, in factor analysis non-orthogonal transformation is conventionally known as *oblique rotation*.

where $f(\mathbf{T}^{(k)}) \equiv \frac{1}{2}f^{(6)}(\mathbf{T}^{(k)})$ and its gradient is

$$\nabla f(\mathbf{T}^{(k)}) = \mathbf{B}^T(\mathbf{Q}^{(k)} \circ \mathbf{Q}^{(k)} \circ \mathbf{Q}^{(k)}).$$

Since $f(\mathbf{T}) \geq \tilde{f}(\mathbf{T})$ for all orthogonal \mathbf{T} (see [84] for a proof), the value of f is increased in each iteration by finding a point $\mathbf{T}^{(k+1)}$ that increases the value of \tilde{f} , which is maximized by $\mathbf{T}^{(k+1)} \equiv \mathbf{U}\mathbf{V}^T$, where $\nabla f(\mathbf{T}) = \mathbf{U}\Sigma\mathbf{V}^T$ is the gradient's SVD [69, pp. 327–328]. Starting with any orthogonal $\mathbf{T}^{(0)}$ (e.g., $\mathbf{T}^{(0)} \equiv \mathbf{I}$) and repeating until convergence, the simple *basic singular value* (BSV) algorithm is:

$$\begin{aligned} \mathbf{Q}^{(k)} &\equiv \mathbf{B}\mathbf{T}^{(k)} \\ \mathbf{U}\Sigma\mathbf{V}^T &\equiv \mathbf{B}^T(\mathbf{Q}^{(k)} \circ \mathbf{Q}^{(k)} \circ \mathbf{Q}^{(k)}) \\ \mathbf{T}^{(k+1)} &\equiv \mathbf{U}\mathbf{V}^T. \end{aligned}$$

Convergence is determined by monitoring $\text{tr}(\Sigma)$, which increases monotonically.

Quartimax solutions approximate orthogonal AJD solutions by minimizing an upper bound on (5.14), since, for orthogonal \mathbf{T} , $\mathbf{Q}^+ = \mathbf{Q}^T$ and

$$\begin{aligned} f^{(3)} &= \sum_{k=1}^n \|\text{off}[\mathbf{Q}^T \text{diag}(\widehat{\mathbf{U}}_k^{(\mathcal{S})})\mathbf{Q}]\|_F^2 = \sum_{k=1}^n \sum_{i \neq j} \langle \mathbf{Q}_i \circ \mathbf{Q}_j, \widehat{\mathbf{U}}_k^{(\mathcal{S})} \rangle^2 \\ &\leq \sum_{k=1}^n \sum_{i \neq j} \|\widehat{\mathbf{U}}_k^{(\mathcal{S})}\|_F^2 \|\mathbf{Q}_i \circ \mathbf{Q}_j\|_F^2 = \sum_{k=1}^n \sigma^{(\mathcal{S})2} \sum_{i \neq j} \|\mathbf{Q}_i \circ \mathbf{Q}_j\|_F^2. \end{aligned}$$

Maximizing $\|\mathbf{Q} \circ \mathbf{Q}\|_F^2$ yields an approximate, albeit non-optimal, joint diagonalization by implicitly minimizing $\sum_{i \neq j} \|\mathbf{Q}_i \circ \mathbf{Q}_j\|_F^2$, the quartimin objective function.

Worth mentioning is the *promax* method, an heuristic which finds oblique, approximate quartimin solutions very quickly by raising the absolute value of a transformed orthogonal basis (found, for example, by the quartimax method) to some power, effectively exaggerating its sharpness. For our needs, the fourth power works well in general. Projection onto \mathbf{B} and a normalization then yields $\mathbf{T}^{(7)}$:

$$\mathbf{Q} \equiv \mathbf{B}\mathbf{T}^{(6)}, \quad \mathbf{T} \equiv \mathbf{B}^T(\mathbf{Q} \circ \mathbf{Q} \circ \mathbf{Q} \circ \mathbf{Q}), \quad \mathbf{T}^{(7)} \equiv \mathbf{T} \text{diag}(\mathbf{T}^{-1}\mathbf{B}^T\mathbf{1}).$$

Table 5.1: Optimization problems

method	objective	constraint
min. spectral error	$\min_{\mathbf{T}} f^{(1)} \equiv \ \Delta \sqrt{\mathbf{c}^{(k-1)} \mathbf{c}^{(k-1)\top}}\ _F^2$	none
min. residual Δ	$\min_{\mathbf{T}} f^{(2)} \equiv \ \Delta\ _F^2$	none
AJD	$\min_{\mathbf{T}} f^{(3)} \equiv \sum_{i=1}^n \left\ \text{off}[\mathbf{Q}^+ \text{diag}(\widehat{\mathbf{U}}_i^{(\mathcal{S})}) \mathbf{Q}] \right\ _F^2$	$\mathbf{T}^\top \mathbf{T} = \mathbf{I}$ or $\text{diag}(\mathbf{T}^\top \mathbf{T}) = \mathbf{1}$
Drew & Finlayson	$\min_{\mathbf{T}} f^{(4)} \equiv \sum_{i=1}^m \ \mathbf{B}(\mathcal{L}^{(i)})(\mathbf{T}^{-\top})_i\ _2^2$	$\text{diag}(\mathbf{T}^\top \mathbf{T})^{-1} = \mathbf{1}$
quartimin	$\min_{\mathbf{T}} f^{(5)} \equiv \sum_{i \neq j} \ \mathbf{Q}_i \circ \mathbf{Q}_j\ _F^2$	$\text{diag}(\mathbf{T}^\top \mathbf{T}) = \mathbf{1}$
quartimax	$\max_{\mathbf{T}} f^{(6)} \equiv \ \mathbf{Q} \circ \mathbf{Q}\ _F^2$	$\mathbf{T}^\top \mathbf{T} = \mathbf{I}$

5.5 By minimizing the total error, $\|\Delta \mathbf{c}\|_F^2$

By taking into account \mathbf{B} , \mathcal{S} and \mathcal{E} , a sharp basis $\mathbf{Q}^{(1)} \equiv \mathbf{B}\mathbf{T}^{(1)}$, found by explicit minimization of $f^{(1)} \equiv \|\Delta \mathbf{c}\|_F^2$, yields the lowest approximation error. AJD, by contrast, neglects \mathcal{E} , minimizing $\sum_{i \neq j, k} \|\mathbf{Q}_i^{+\top} \circ \mathbf{Q}_j \circ \mathcal{S}_k\|_F^2$, while spectral sharpening neglects both \mathcal{S} and \mathcal{E} when minimizing $\sum_{i \neq j} \|\mathbf{Q}_i^{+\top} \circ \mathbf{Q}_j\|_F^2$, which depends only on \mathbf{B} . Examination of the SVDs of Δ and $\mathbf{c} \equiv \mathbf{c}^{(1)}$ shows that the error incurred by minimizing $f^{(1)}$ is determined by the relative orientation of the right singular vectors of the residual Δ and the left singular vectors of \mathbf{c} (which does not depend on \mathbf{T}), weighted by their respective singular values:

$$\|\Delta \mathbf{c}\| = \|\mathbf{U}^{(\Delta)} \Sigma^{(\Delta)} \mathbf{V}^{(\Delta)\top} \mathbf{U}^{(\mathbf{c})} \Sigma^{(\mathbf{c})} \mathbf{V}^{(\mathbf{c})\top}\| = \|\Sigma^{(\Delta)} \mathbf{V}^{(\Delta)\top} \mathbf{U}^{(\mathbf{c})} \Sigma^{(\mathbf{c})}\|.$$

As the value of m increases, the smallest singular values of \mathbf{c} decrease rapidly towards zero. This creates additional degrees of freedom for an optimization algorithm to orient the leading vectors of $\mathbf{V}^{(\Delta)}$ in directions approximately orthogonal to those of the leading vectors of $\mathbf{U}^{(\mathbf{c})}$, while increasing $\|\Delta\| = \|\Sigma^{(\Delta)}\|$ as $\|\Delta \mathbf{c}\|$ decreases.

When a gradient-based method is used to minimize $f^{(1)}$, explicit and efficient computation of the gradient $\nabla f^{(1)}$ (instead of, for example, its approximation by finite differences) accelerates convergence by requiring fewer iterations and computing

them more quickly. The gradient of $f \equiv f^{(1)}(\mathbf{T}, \mathbf{B}, \widehat{\mathbf{U}}^{(\mathcal{S})}, \mathbf{c} \equiv \mathbf{c}^{(1)})$,

$$\nabla f = \frac{\partial}{\partial \mathbf{T}} \|\Delta \sqrt{\mathbf{c}\mathbf{c}^T}\|_F^2 = \frac{\partial}{\partial \mathbf{T}} \text{tr}(\Delta \mathbf{c}\mathbf{c}^T \Delta^T),$$

may be computed at a cost of $O(n \cdot m^5)$ one partial derivative entry at a time as

$$\nabla f_{ab} = 2 \sum_{i=1}^n \text{tr}(\Delta^{(i)} \mathbf{c}\mathbf{c}^T \frac{\partial}{\partial \mathbf{T}_{ab}} \Delta^{(i)T}),$$

but a more efficient approach, at a cost of $O(n \cdot m^3)$, expresses the gradient as the sum $\sum_{i=1}^n \nabla f^{(i)}$, each term of which is a single matrix computation for one column of $\widehat{\mathbf{U}}^{(\mathcal{S})}$, derived using the matrix differential (see [112]),

$$\begin{aligned} \mathbf{d}f^{(i)} &= 2 \sum_{i=1}^n \text{tr}(\Delta^{(i)} \mathbf{c}\mathbf{c}^T \mathbf{d}\Delta^{(i)T}), \text{ where} \\ \mathbf{d}\Delta^{(i)} &= \mathbf{T}[-\mathbf{T}^{-1} \mathbf{d}\mathbf{T} \text{diag}(\widetilde{\mathbf{u}}_i^{(\mathcal{S})}) + \text{diag}(\mathbf{T}^{-1} \mathbf{d}\mathbf{T} \widetilde{\mathbf{u}}_i^{(\mathcal{S})}) + \text{diag}(\widetilde{\mathbf{u}}_i^{(\mathcal{S})}) \mathbf{T}^{-1} \mathbf{d}\mathbf{T}] \mathbf{T}^{-1} \end{aligned}$$

Rearrangement and the following facts, for square matrices \mathbf{A} and \mathbf{B} and vector \mathbf{c} ,

$$\begin{aligned} \text{tr}(\mathbf{A} + \mathbf{B}) &= \text{tr}(\mathbf{A}) + \text{tr}(\mathbf{B}) \\ \text{tr}(\mathbf{A}\mathbf{B}) &= \text{tr}(\mathbf{B}\mathbf{A}) = \text{tr}(\mathbf{A}^T \mathbf{B}^T) \\ \text{tr}[\mathbf{B} \text{diag}(\mathbf{A}\mathbf{c})] &= \text{tr}[\mathbf{A}\mathbf{c}\mathbf{1}^T \text{ddiag}(\mathbf{B})], \end{aligned}$$

yield $\text{tr}(\Delta^{(i)} \mathbf{c}\mathbf{c}^T \mathbf{d}\Delta^{(i)T}) \equiv \text{tr}(\mathbf{X}^{(i)T} \mathbf{d}\mathbf{T})$, where

$$\begin{aligned} \mathbf{X}^{(i)} &\equiv \mathbf{T}^{-T}[-\mathbf{Y}^{(i)} \text{diag}(\widetilde{\mathbf{u}}_i^{(\mathcal{S})}) + \text{diag}(\widetilde{\mathbf{u}}_i^{(\mathcal{S})}) \mathbf{Y}^{(i)} + \text{ddiag}(\mathbf{Y}^{(i)}) \mathbf{1}\widetilde{\mathbf{u}}_i^{(\mathcal{S})T}] \text{ and} \\ \mathbf{Y}^{(i)} &\equiv \mathbf{T}^T \Delta^{(i)} \mathbf{c}\mathbf{c}^T \mathbf{T}^{-T}. \end{aligned}$$

Thus, $\nabla f = 2 \sum_{i=1}^n \mathbf{X}^{(i)}$.

Iterative methods benefit by starting at a point $\mathbf{T}^{(0)}$ that is nearer to a solution. Found quickly, quartimax solutions are good candidates for starting points when minimizing $f^{(1)}$, although AJD solutions, which take more time to find, will likely be closer to the desired solutions, which, as shown in Table 5.2, is true in most cases considered here. For example, a quartimax solution for the 9-dimensional $\mathbf{B}^{(2)}$ basis for Munsell $\mathcal{C}^{(2)}$ SPDs formed with illuminants D65, A and F2, can be found in

40 microseconds⁶ after three iterations of the Neuhaus and Wrigley pairwise vector method with the initialization $\mathbf{T}^{(0)} \equiv \mathbf{I}$ and a convergence criterion of a relative change of less than 10^{-3} in the objective function's value. The UWAJD algorithm, starting with the quartimax solution and using the same stopping criterion, converges in three iterations requiring 80 additional microseconds. The added time is more than compensated for when minimizing $f^{(1)}$ with the BFGS algorithm. With the same stopping criterion, setting $\mathbf{T}^{(0)}$ to the UWAJD solution instead of the quartimax solution reduces convergence time by 1.8 ms, or 42%, as shown in Table 5.2. Also shown in Table 5.2 is that the BFGS algorithm converges more quickly when initialized with any approximate solution than it does, on average, when $\mathbf{T}^{(0)}$ is a normalized, random orthogonal matrix. However, because a nonlinear optimization algorithm can converge to a local, non-optimal minimum, a better solution can often be found by selecting the best one obtained from a collection of random starting points.

Figs. 5.20-5.25 show plots of the first nine vectors of orthonormal bases $\mathbf{B}^{(1)}$ for multi-illuminant $\mathcal{C}^{(1)}$ SPDs constructed with Munsell, Natural Colors and METACOW reflectance spectra and illuminants $\mathcal{E} \equiv \{\text{D65, A, F2}\}$, along with their transformations as sharp bases $\mathbf{Q}^{(1)} \equiv \mathbf{B}^{(1)}\mathbf{T}^{(1)}$ found by minimizing $f^{(1)}$. Figs. 5.26-5.49 show the same plots for single- and multi-illuminant $\mathcal{C}^{(2)}$ SPDs. Also included, for comparison, are maximally sharp bases $\mathbf{Q}^{(5)} \equiv \mathbf{B}^{(1)}\mathbf{T}^{(5)}$ found by minimizing $f^{(5)}$. As is the case for the sharpened single-illuminant, single-reflection bases shown in Figs. 5.1-5.18, the values of $f^{(5)}$ displayed in the rightmost columns of plots show that $\mathbf{Q}^{(1)}$ bases are considerably less sharp than the corresponding $\mathbf{Q}^{(5)}$ bases, although, as will be shown in the next chapter, they yield lower approximation errors.

⁶All performance results reported here and elsewhere were obtained on the following platform, using double-precision floating point arithmetic on the CPU or single-precision on the GPU: Intel Core i7-3930K CPU, Ubuntu 16.04.2 LTS (AMD64, Linux kernel 4.4.0-109-generic), Nvidia GeForce GTX Titan Black GPU (driver version 384.111 with OpenGL 4.5.0). All code was written in C++ and compiled with g++ version 5.4.0 of the GNU Compiler Collection.

Table 5.2: BFGS $f^{(1)}$ call counts and run times

		source of BFGS initial point $\mathbf{T}^{(0)}$					
		UWAJD		quartimax		random	
		$f^{(1)}$	time	$f^{(1)}$	time	$f^{(1)}$	time
m		calls	(ms)	calls	(ms)	calls	(ms)
Munsell							
4		13	0.34	20	0.35	39	0.68
5		19	0.55	27	0.64	74	1.6
6		28	0.85	28	0.79	102	2.4
7		45	1.7	45	1.7	158	5.2
8		41	2.2	73	3.3	195	8.1
9		40	2.5	76	4.3	220	11
Natural Colors							
4		32	0.73	36	0.68	56	1.1
5		28	0.87	28	0.82	103	2.8
6		38	1.3	45	1.4	150	4.5
7		72	3.3	53	2.4	182	7.8
8		55	3.5	53	3.2	211	12
9		51	3.8	114	7.8	229	15
METACOW							
4		11	0.23	16	0.24	49	0.68
5		27	0.61	37	0.72	83	1.6
6		24	0.65	60	1.3	121	2.6
7		48	1.6	70	2.2	168	5.1
8		49	2.2	54	2.3	191	7.5
9		93	4.9	101	5.1	215	11

For dimensionalities $m \equiv 4, \dots, 9$, the objective function call counts and run times (in milliseconds, rounded to two significant figures) for different initializations of the BFGS minimization of $f^{(1)}$ applied to $\mathcal{C}^{(2)}$ SPDs constructed with illuminants D65, A and F2. Objective functions were constructed using $\{\hat{\mathbf{U}}_i^{(\mathcal{S})}\}_{i=1}^d$, the leading d left singular vectors of \mathcal{S} , weighted by their respective singular values, where d was chosen to account for 99.99% of the sum of the squared singular values (or, that is, the variance about $\mathbf{0}$). For Munsell, Natural Colors and METACOW reflectance spectra, $d \equiv 10, 18$ and 10 , respectively. All stopping criteria were a relative change of less than 10^{-3} in the value of $f^{(1)}$. UWAJD iterations were initialized by quartimax solutions, and quartimax iterations were initialized by the identity matrix. Random matrices were obtained by normalization of the orthogonal factors of the QR decompositions of $m \times m$ randomized matrices. All call counts and run times, including $\mathbf{T}^{(0)}$ initializations, are averages of 1,000 runs.

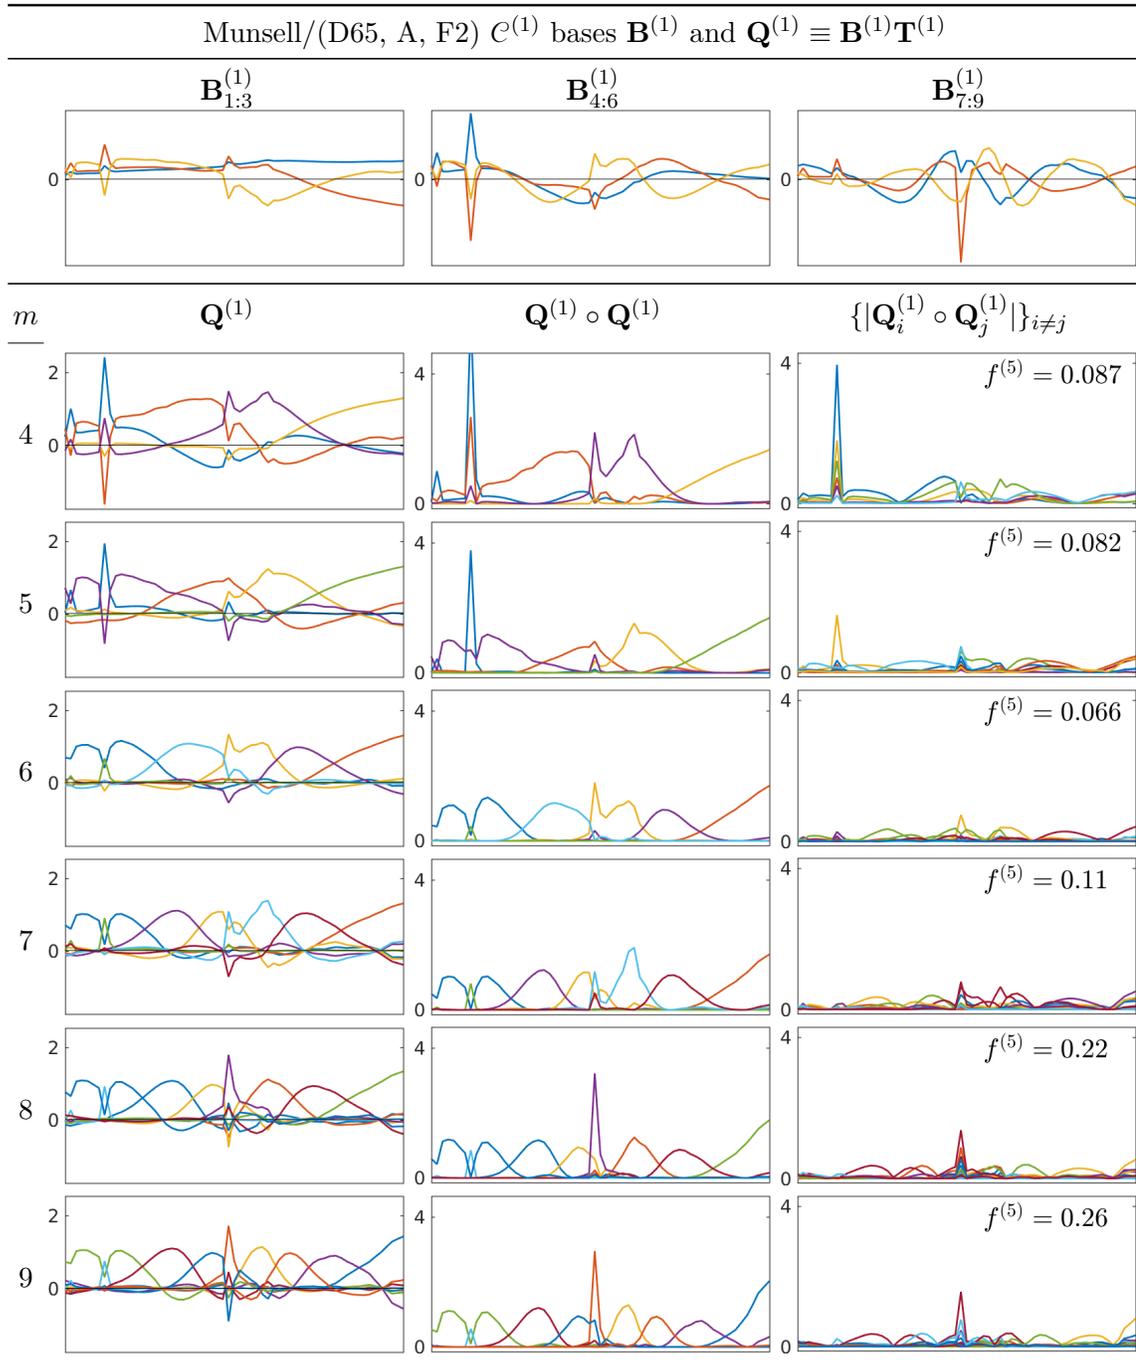


Figure 5.20: For Munsell $\mathcal{C}^{(1)}$ color signals formed with D65, A and F2, the orthonormal basis $\mathbf{B}^{(1)}$, the sharp bases $\mathbf{Q}^{(1)} \equiv \mathbf{B}^{(1)}\mathbf{T}^{(1)}$ (for dimensionalities $m = 4, \dots, 9$), their squares and the absolute values of the pairwise products of their columns.

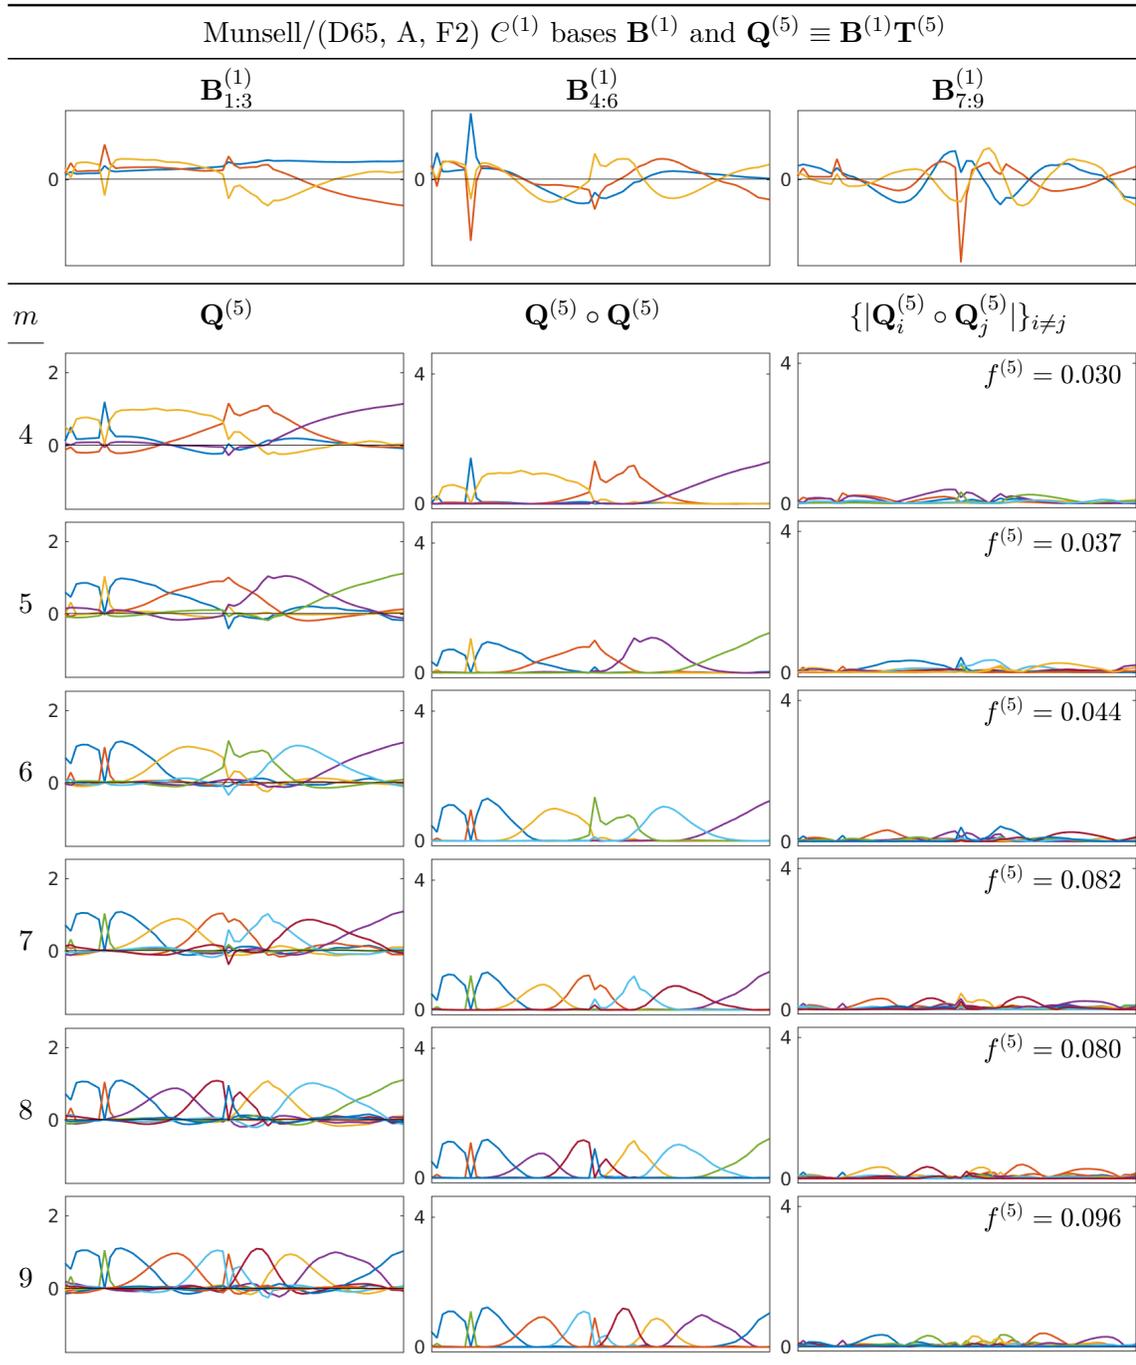


Figure 5.21: For Munsell $\mathcal{C}^{(1)}$ color signals formed with D65, A and F2, the orthonormal basis $\mathbf{B}^{(1)}$, the sharp bases $\mathbf{Q}^{(5)} \equiv \mathbf{B}^{(1)}\mathbf{T}^{(5)}$ (for dimensionalities $m = 4, \dots, 9$), their squares and the absolute values of the pairwise products of their columns.

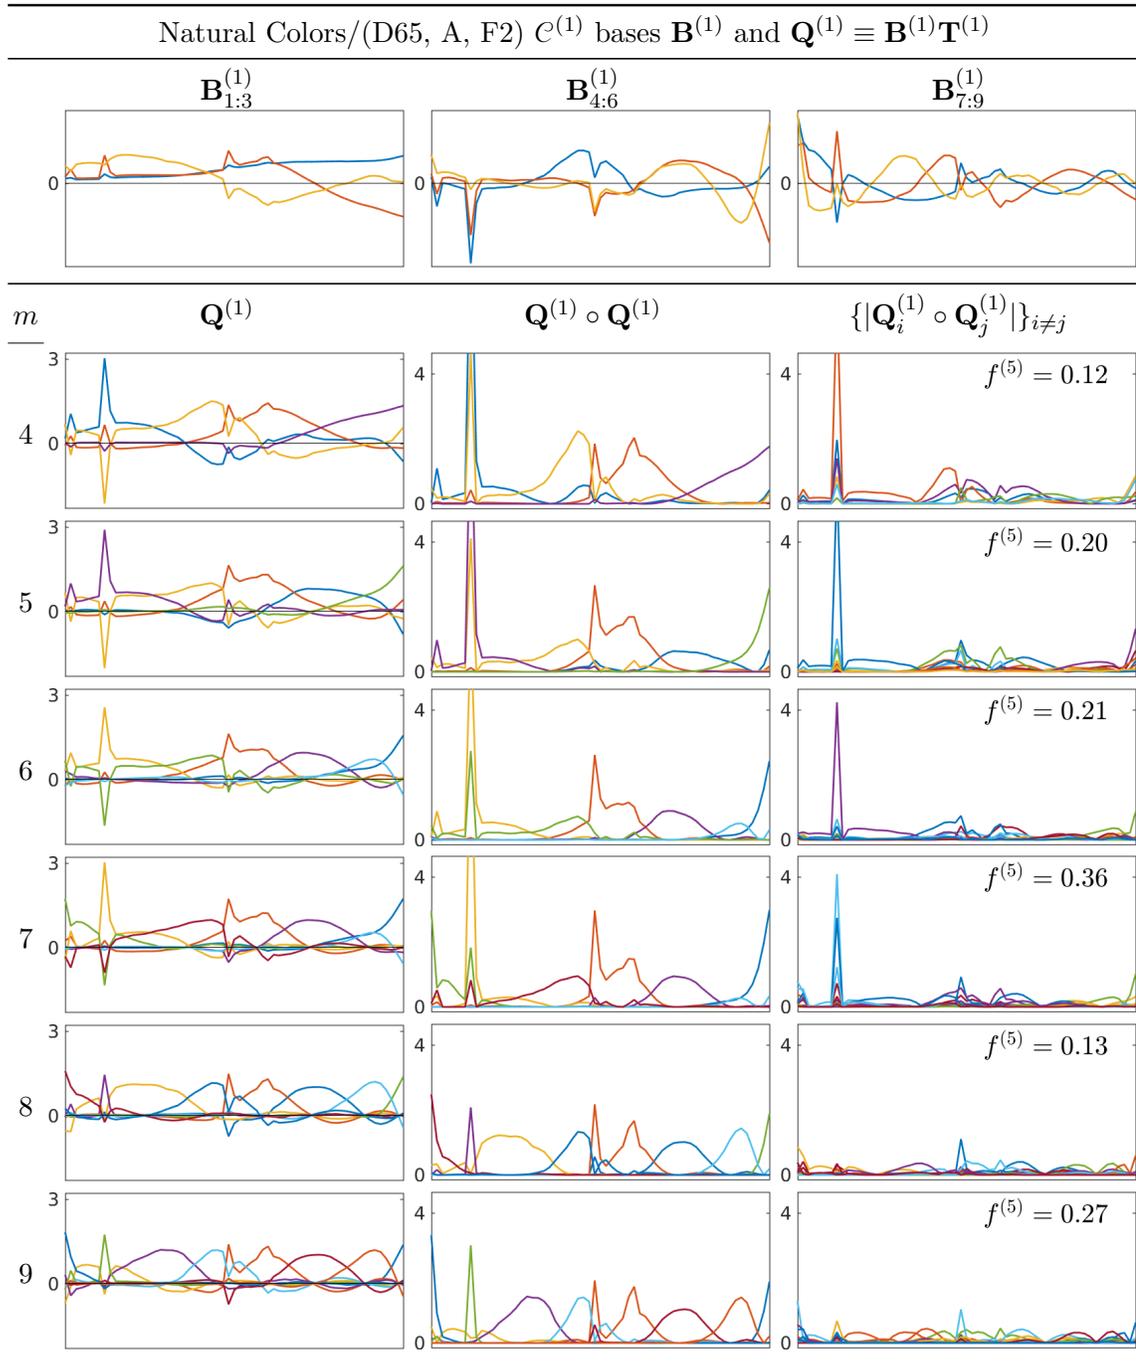


Figure 5.22: For Natural Colors $\mathcal{C}^{(1)}$ color signals formed with D65, A and F2, the orthonormal basis $\mathbf{B}^{(1)}$, the sharp bases $\mathbf{Q}^{(1)} \equiv \mathbf{B}^{(1)}\mathbf{T}^{(1)}$ (for dimensionalities $m = 4, \dots, 9$), their squares and the absolute values of the pairwise products of their columns.

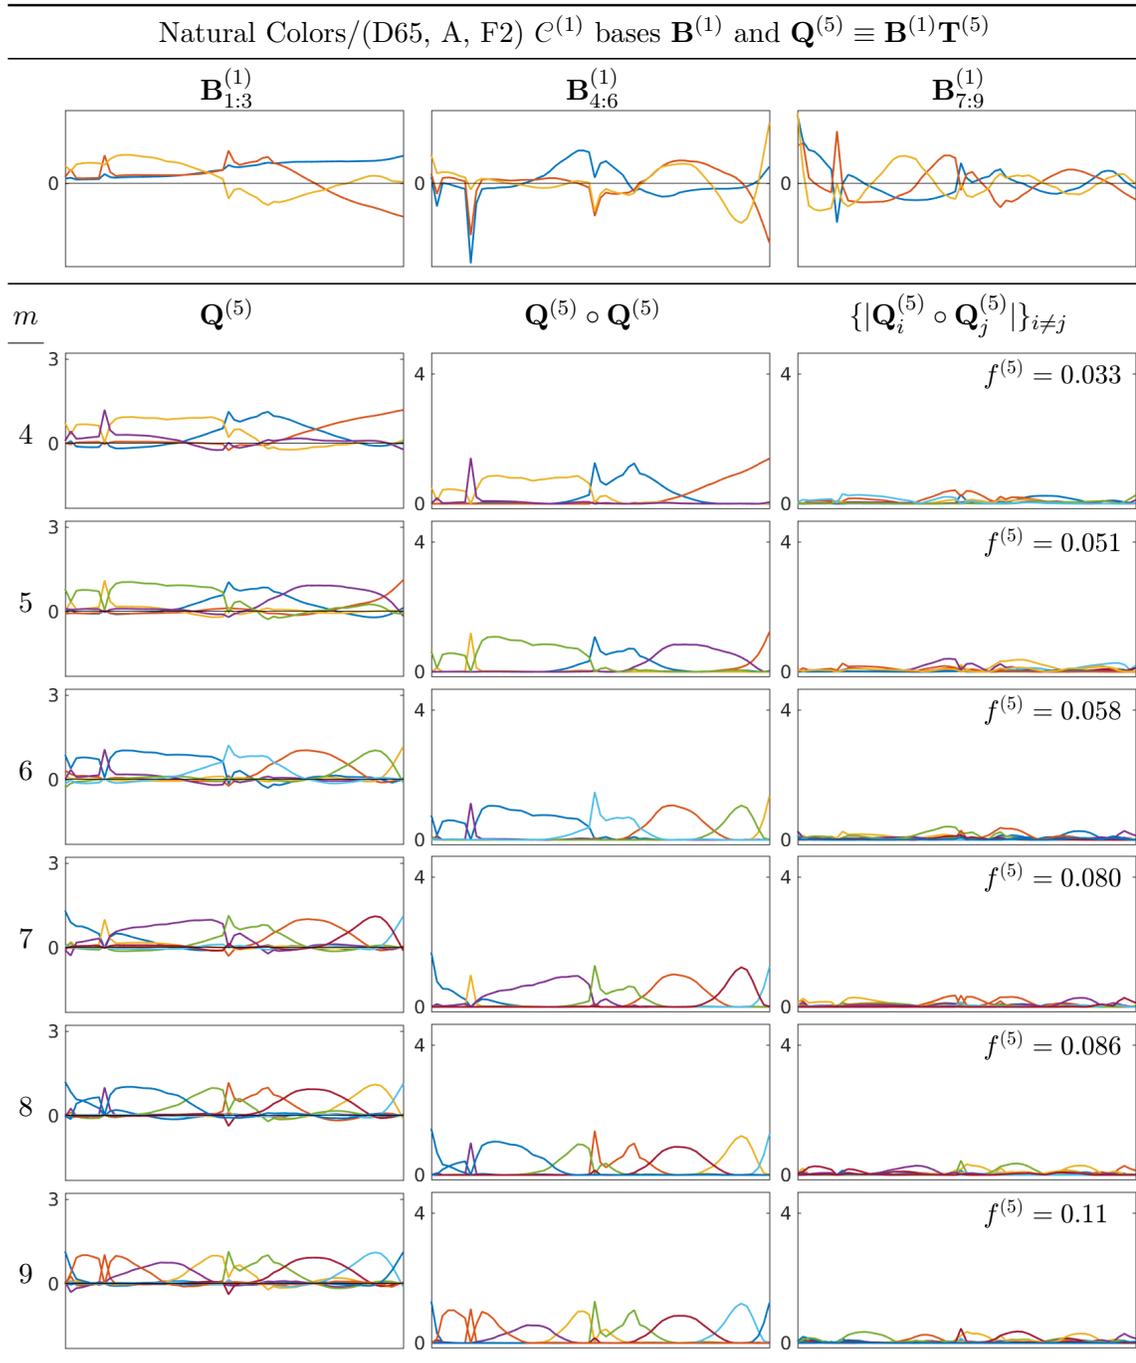


Figure 5.23: For Natural Colors $\mathcal{C}^{(1)}$ color signals formed with D65, A and F2, the orthonormal basis $\mathbf{B}^{(1)}$, the sharp bases $\mathbf{Q}^{(5)} \equiv \mathbf{B}^{(1)}\mathbf{T}^{(5)}$ (for dimensionalities $m = 4, \dots, 9$), their squares and the absolute values of the pairwise products of their columns.

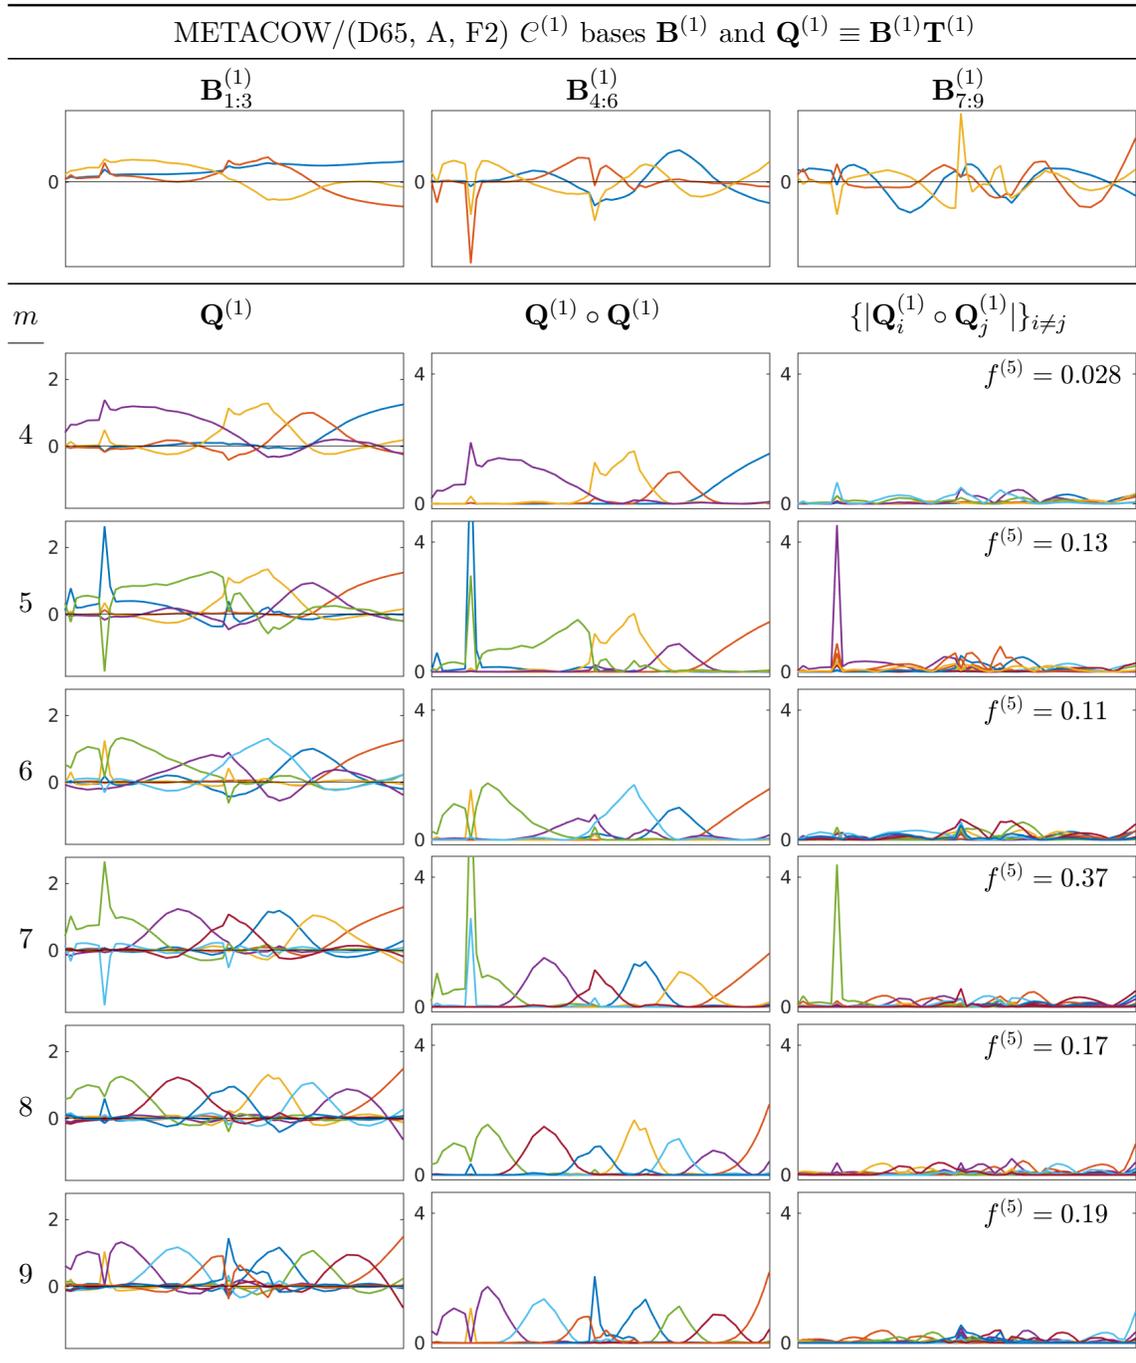


Figure 5.24: For METACOW $\mathcal{C}^{(1)}$ color signals formed with D65, A and F2, the orthonormal basis $\mathbf{B}^{(1)}$, the sharp bases $\mathbf{Q}^{(1)} \equiv \mathbf{B}^{(1)}\mathbf{T}^{(1)}$ (for dimensionalities $m = 4, \dots, 9$), their squares and the absolute values of the pairwise products of their columns.

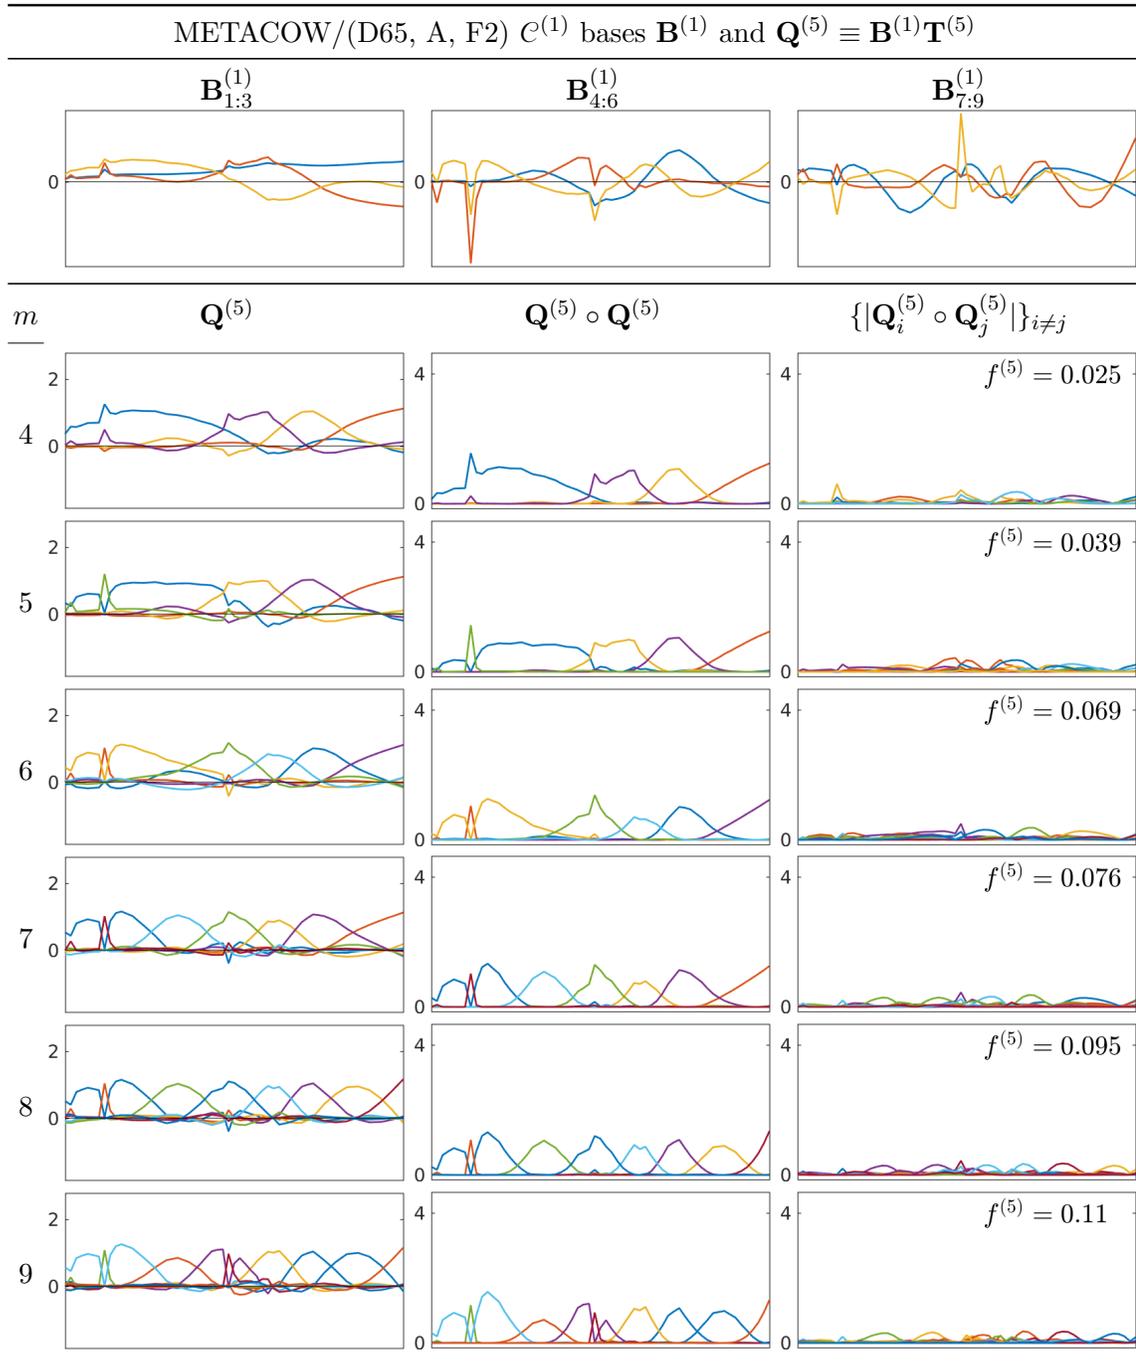


Figure 5.25: For METACOW $\mathcal{C}^{(1)}$ color signals formed with D65, A and F2, the orthonormal basis $\mathbf{B}^{(1)}$, the sharp bases $\mathbf{Q}^{(5)} \equiv \mathbf{B}^{(1)}\mathbf{T}^{(5)}$ (for dimensionalities $m = 4, \dots, 9$), their squares and the absolute values of the pairwise products of their columns.

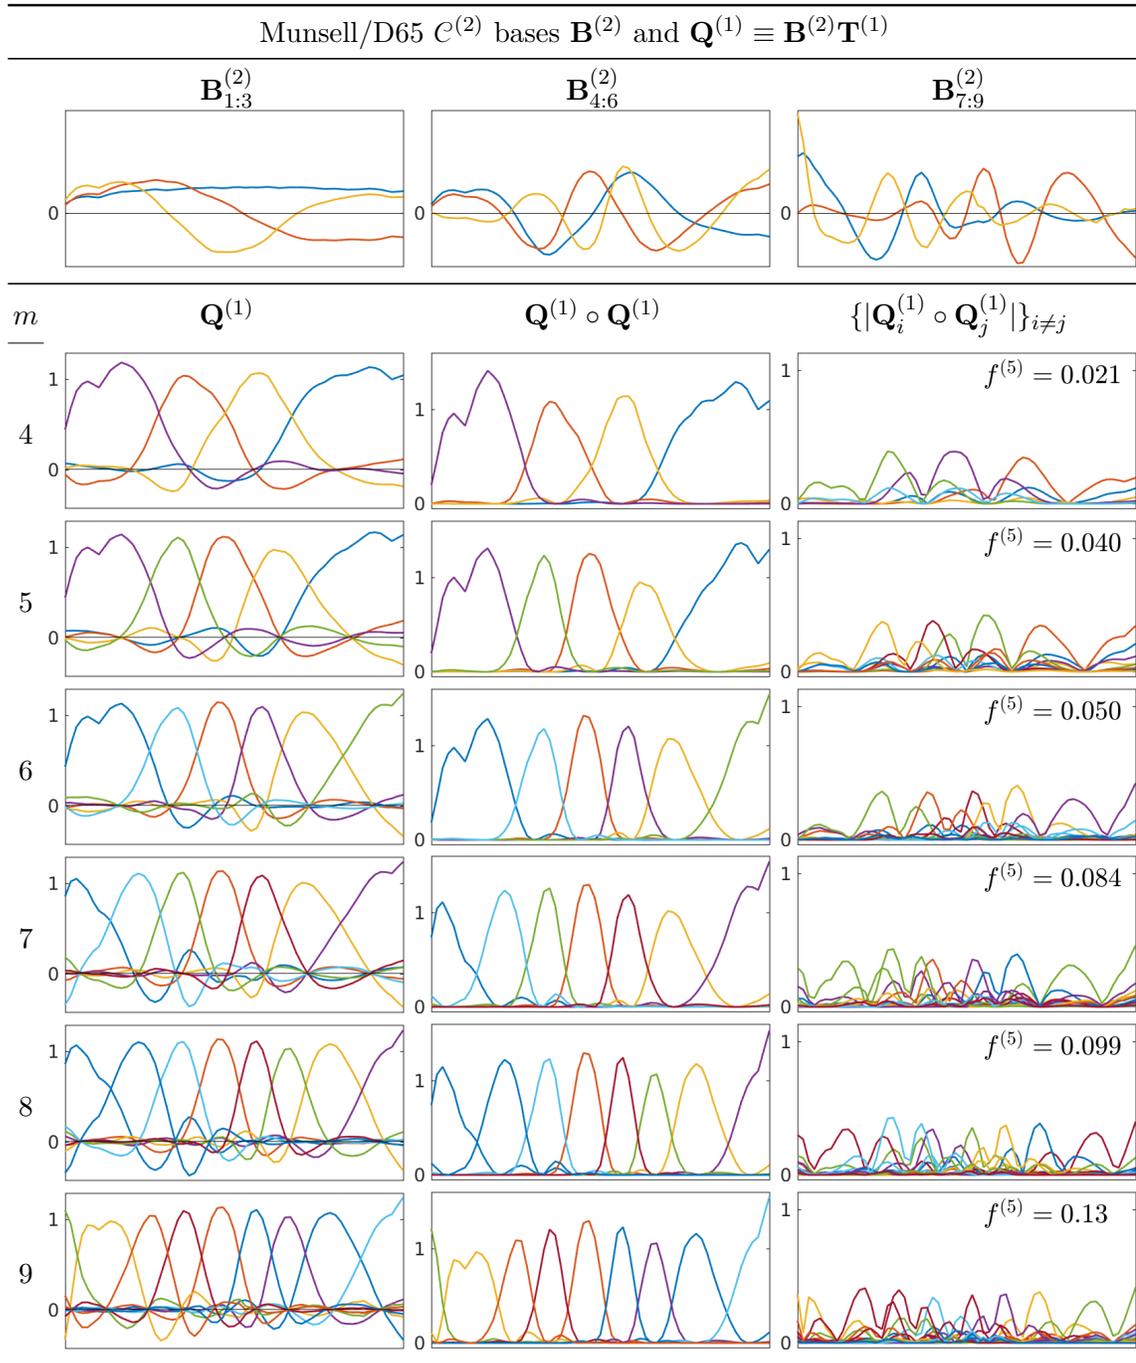


Figure 5.26: For Munsell $C^{(2)}$ color signals formed with D65, the orthonormal basis $\mathbf{B}^{(2)}$, the sharp bases $\mathbf{Q}^{(1)} \equiv \mathbf{B}^{(2)}\mathbf{T}^{(1)}$ (for dimensionalities $m \equiv 4, \dots, 9$), their squares and the absolute values of the pairwise products of their columns.

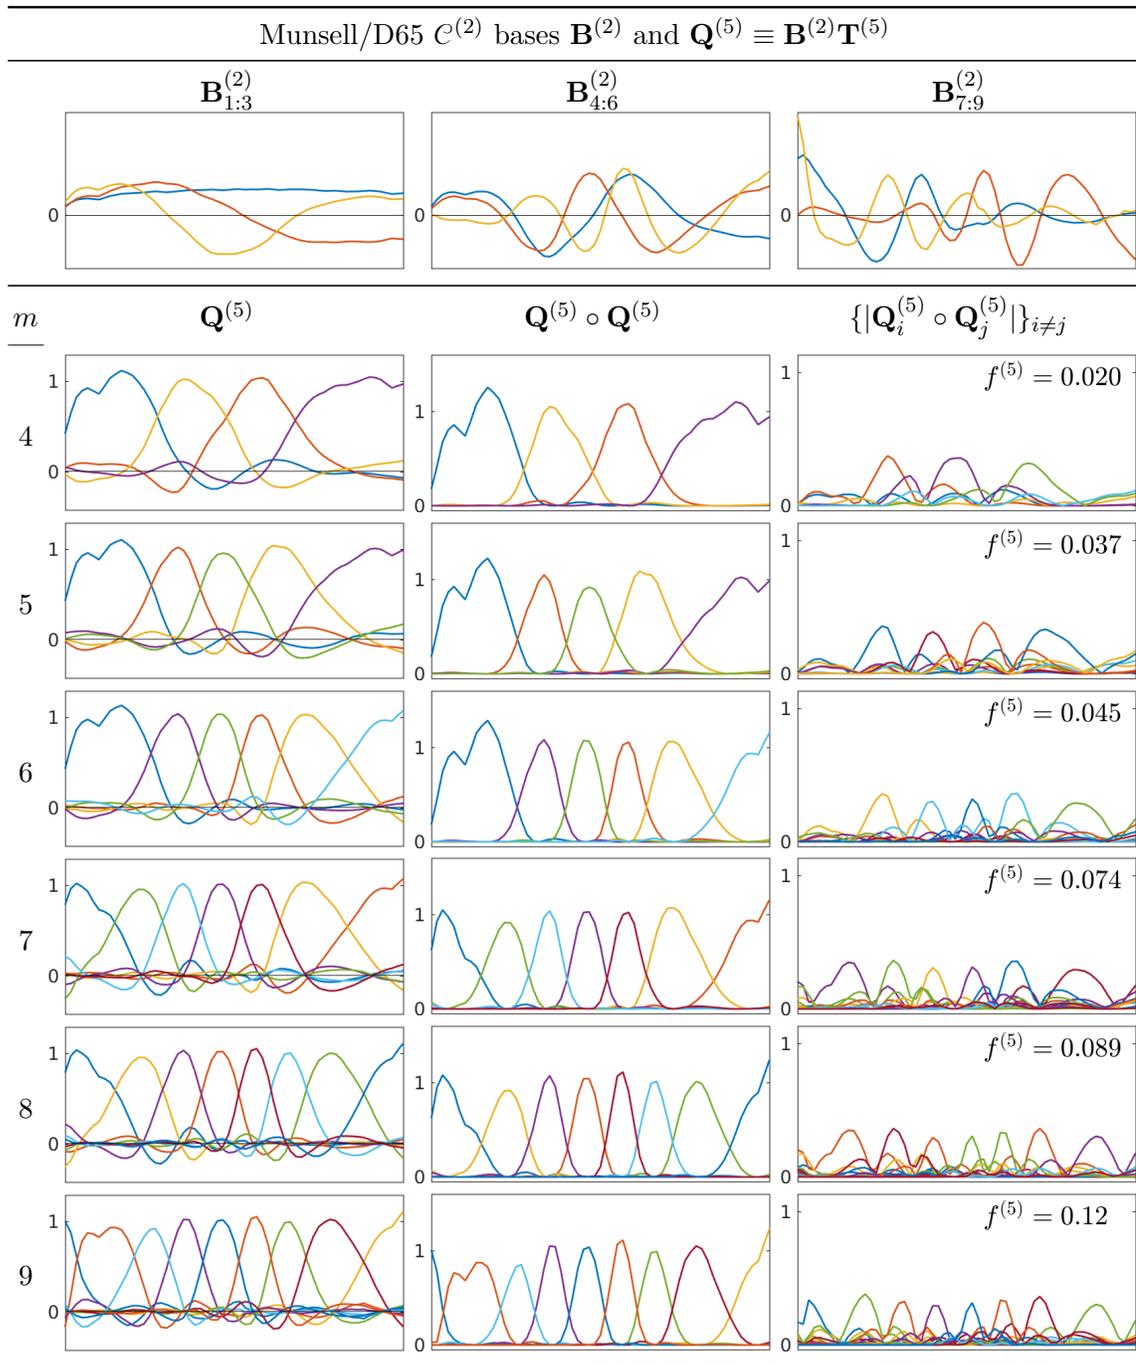


Figure 5.27: For Munsell $C^{(2)}$ color signals formed with D65, the orthonormal basis $\mathbf{B}^{(2)}$, the sharp bases $\mathbf{Q}^{(5)} \equiv \mathbf{B}^{(2)}\mathbf{T}^{(5)}$ (for dimensionalities $m \equiv 4, \dots, 9$), their squares and the absolute values of the pairwise products of their columns.

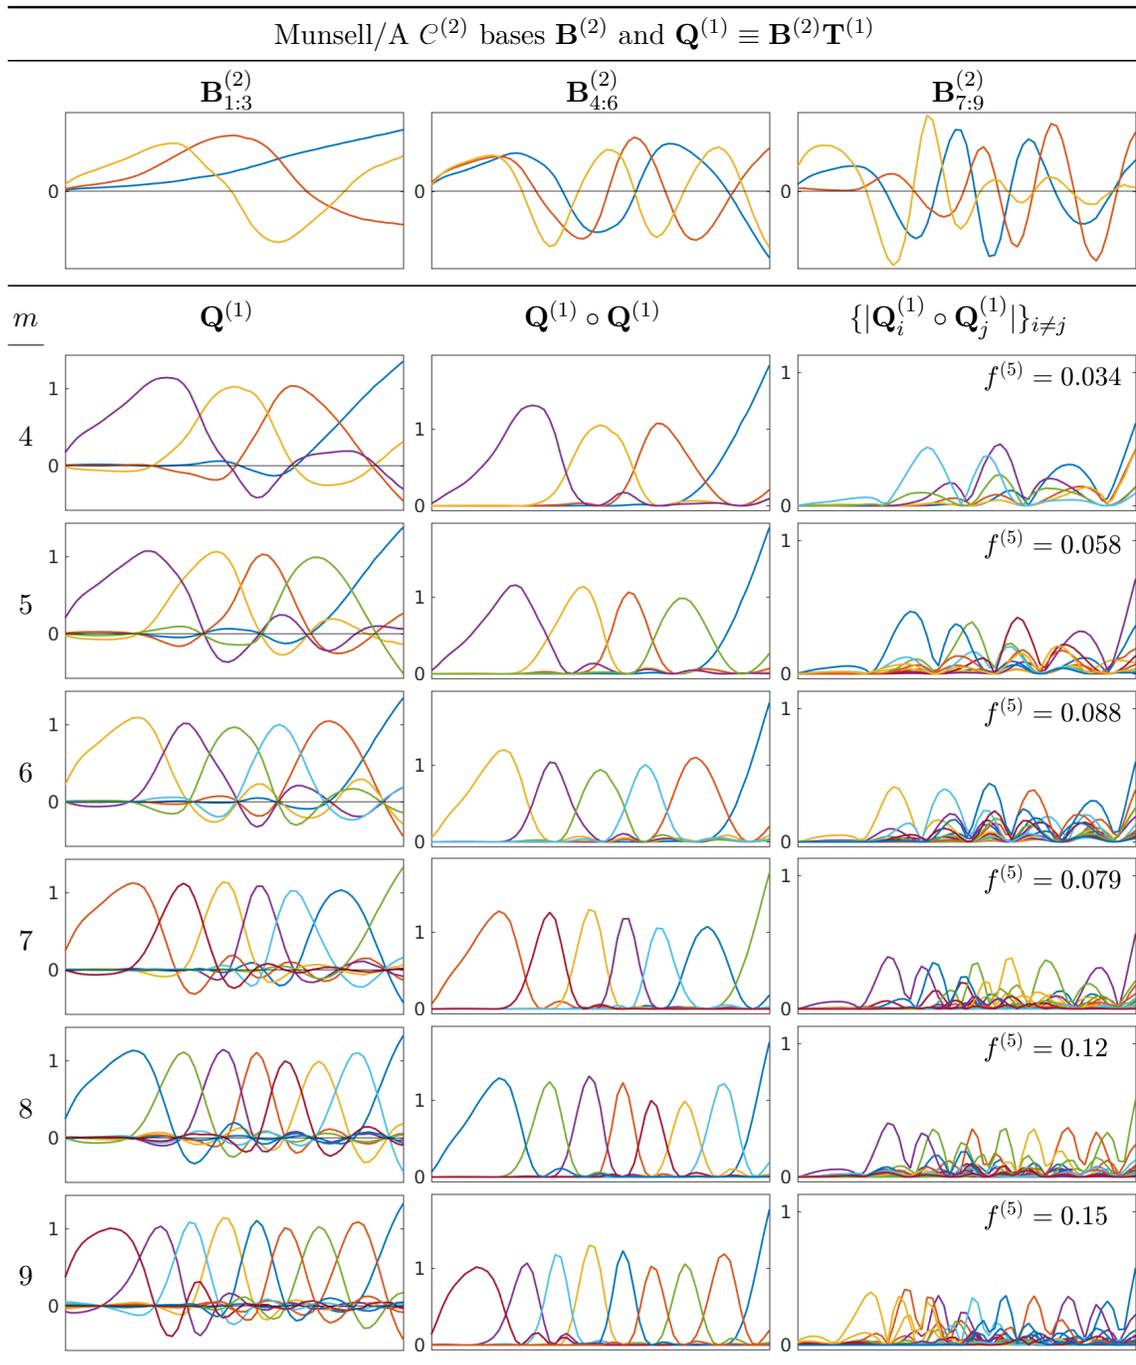


Figure 5.28: For Munsell $C^{(2)}$ color signals formed with illuminant A, the orthonormal basis $\mathbf{B}^{(2)}$, the sharp bases $\mathbf{Q}^{(1)} \equiv \mathbf{B}^{(2)}\mathbf{T}^{(1)}$ (for dimensionalities $m \equiv 4, \dots, 9$), their squares and the absolute values of the pairwise products of their columns.

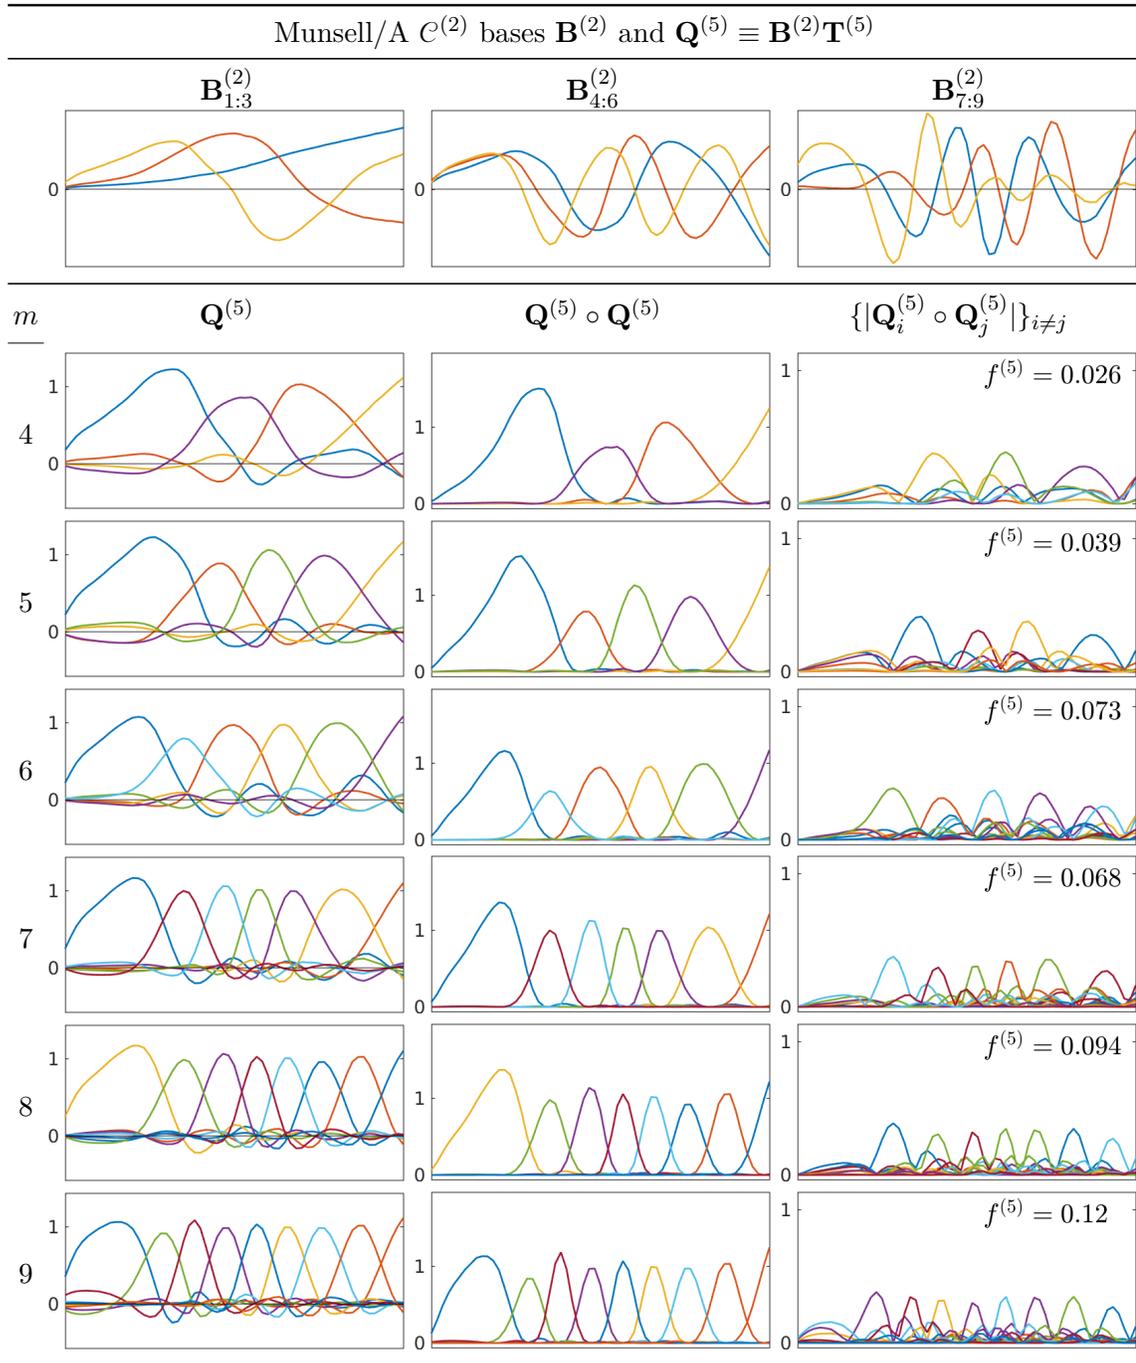


Figure 5.29: For Munsell $C^{(2)}$ color signals formed with illuminant A, the orthonormal basis $\mathbf{B}^{(2)}$, the sharp bases $\mathbf{Q}^{(5)} \equiv \mathbf{B}^{(2)}\mathbf{T}^{(5)}$ (for dimensionalities $m \equiv 4, \dots, 9$), their squares and the absolute values of the pairwise products of their columns.

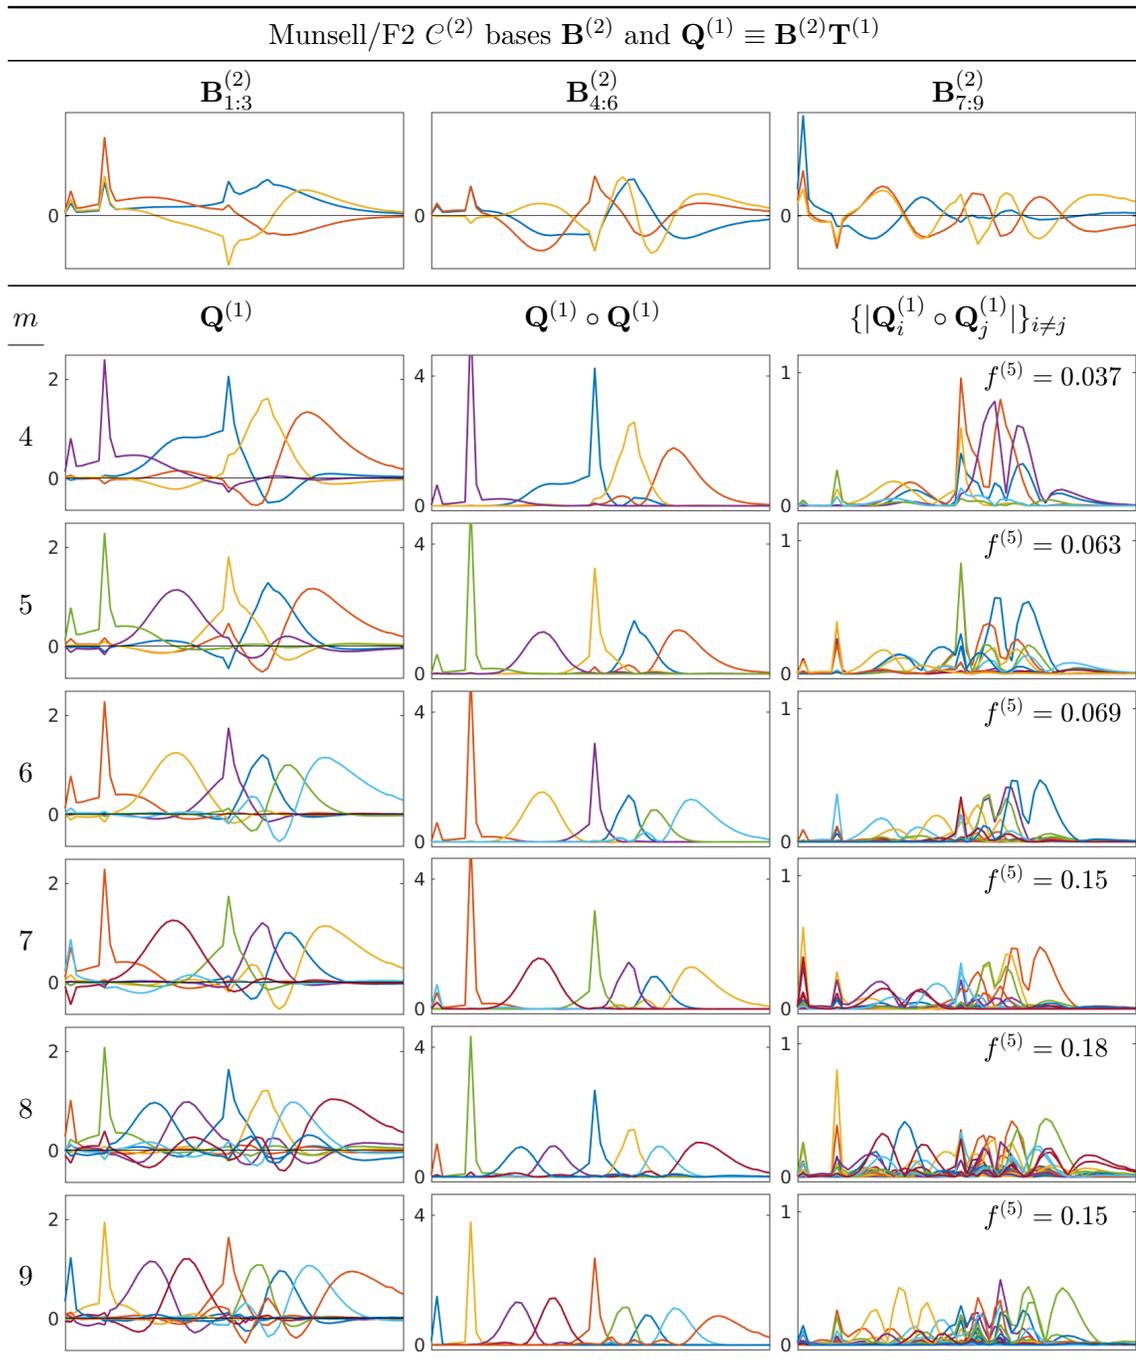


Figure 5.30: For Munsell $\mathcal{C}^{(2)}$ color signals formed with F2, the orthonormal basis $\mathbf{B}^{(2)}$, the sharp bases $\mathbf{Q}^{(1)} \equiv \mathbf{B}^{(2)}\mathbf{T}^{(1)}$ (for dimensionalities $m \equiv 4, \dots, 9$), their squares and the absolute values of the pairwise products of their columns.

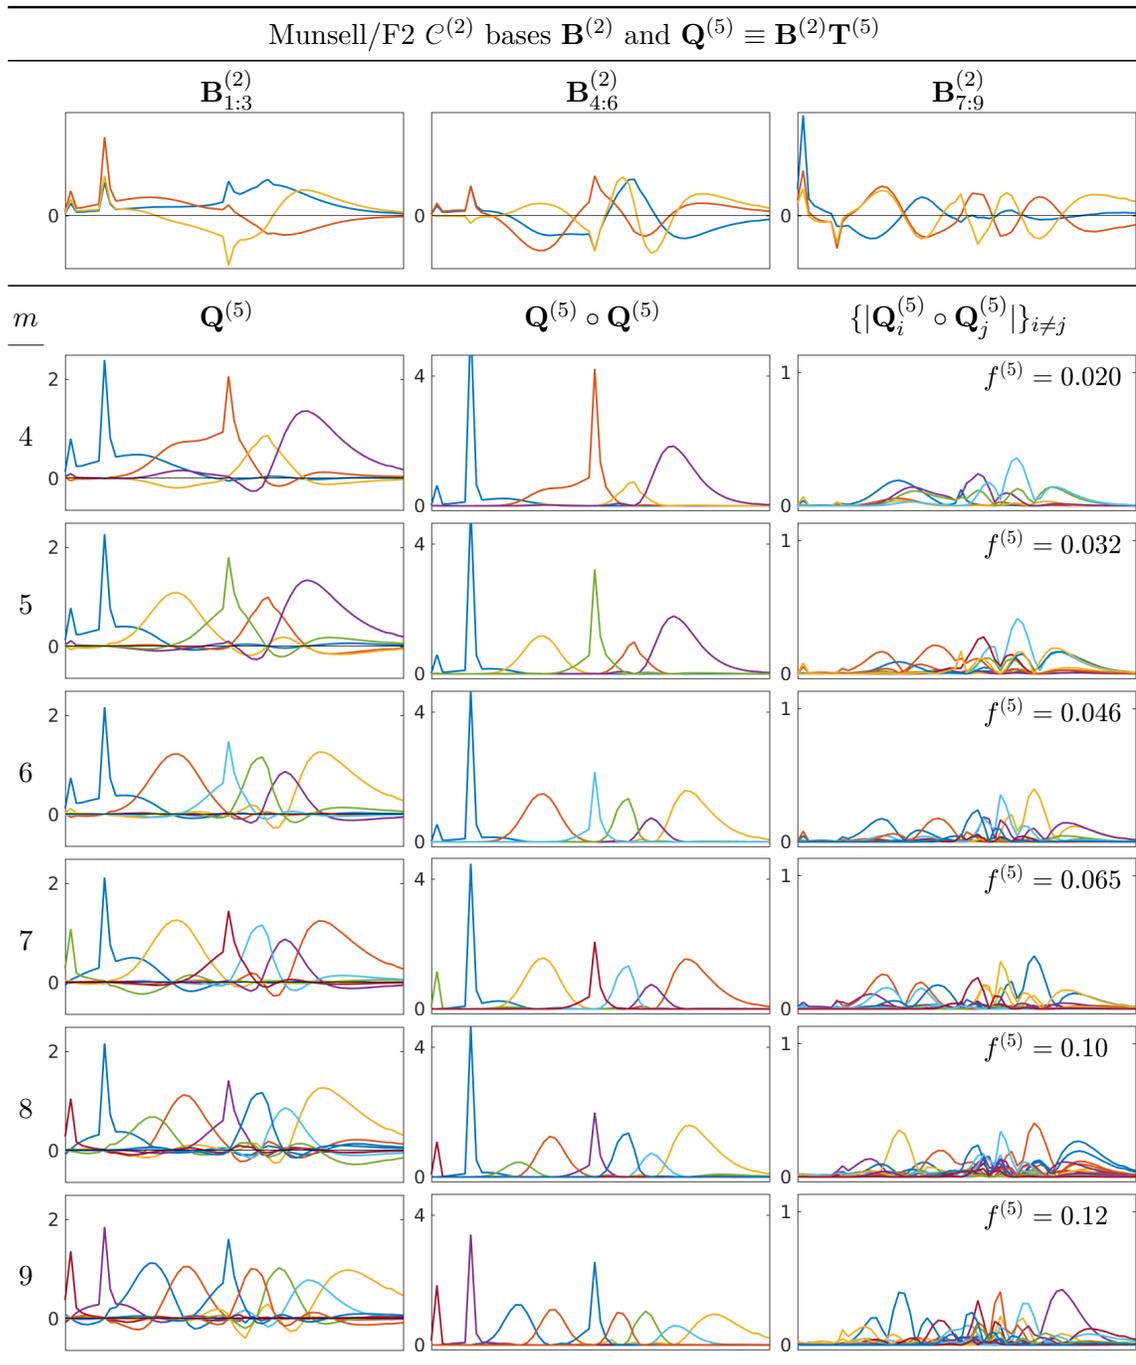


Figure 5.31: For Munsell $\mathcal{C}^{(2)}$ color signals formed with F2, the orthonormal basis $\mathbf{B}^{(2)}$, the sharp bases $\mathbf{Q}^{(5)} \equiv \mathbf{B}^{(2)}\mathbf{T}^{(5)}$ (for dimensionalities $m \equiv 4, \dots, 9$), their squares and the absolute values of the pairwise products of their columns.

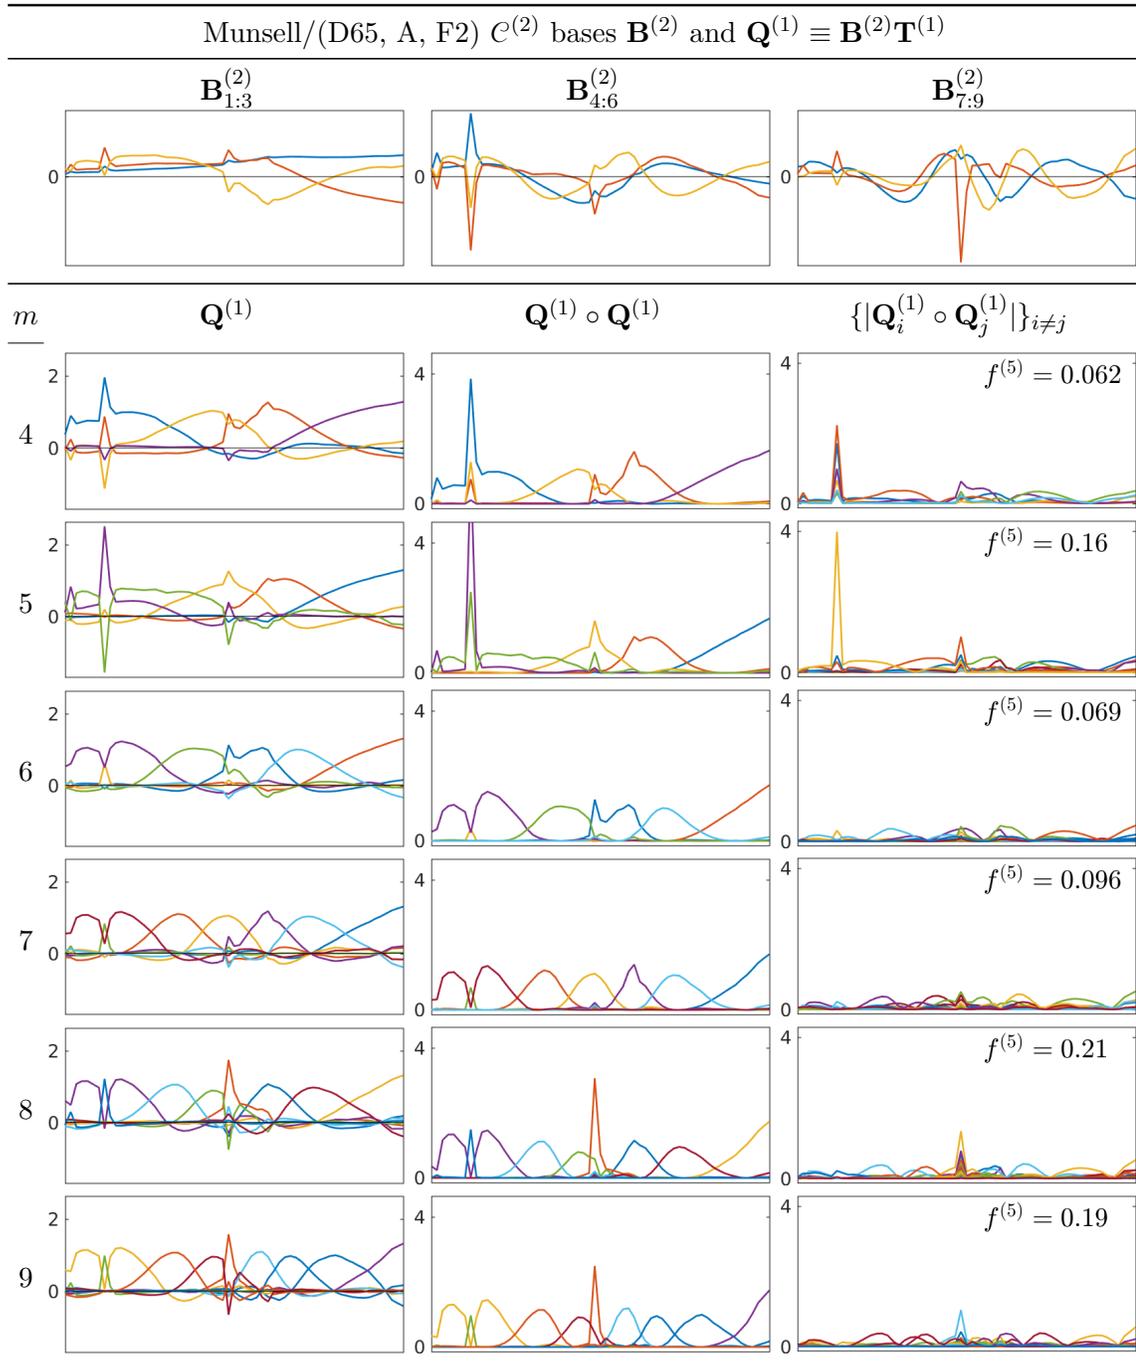


Figure 5.32: For Munsell $\mathcal{C}^{(2)}$ color signals formed with D65, A and F2, the orthonormal basis $\mathbf{B}^{(2)}$, the sharp bases $\mathbf{Q}^{(1)} \equiv \mathbf{B}^{(2)}\mathbf{T}^{(1)}$ (for dimensionalities $m \equiv 4, \dots, 9$), their squares and the absolute values of the pairwise products of their columns.

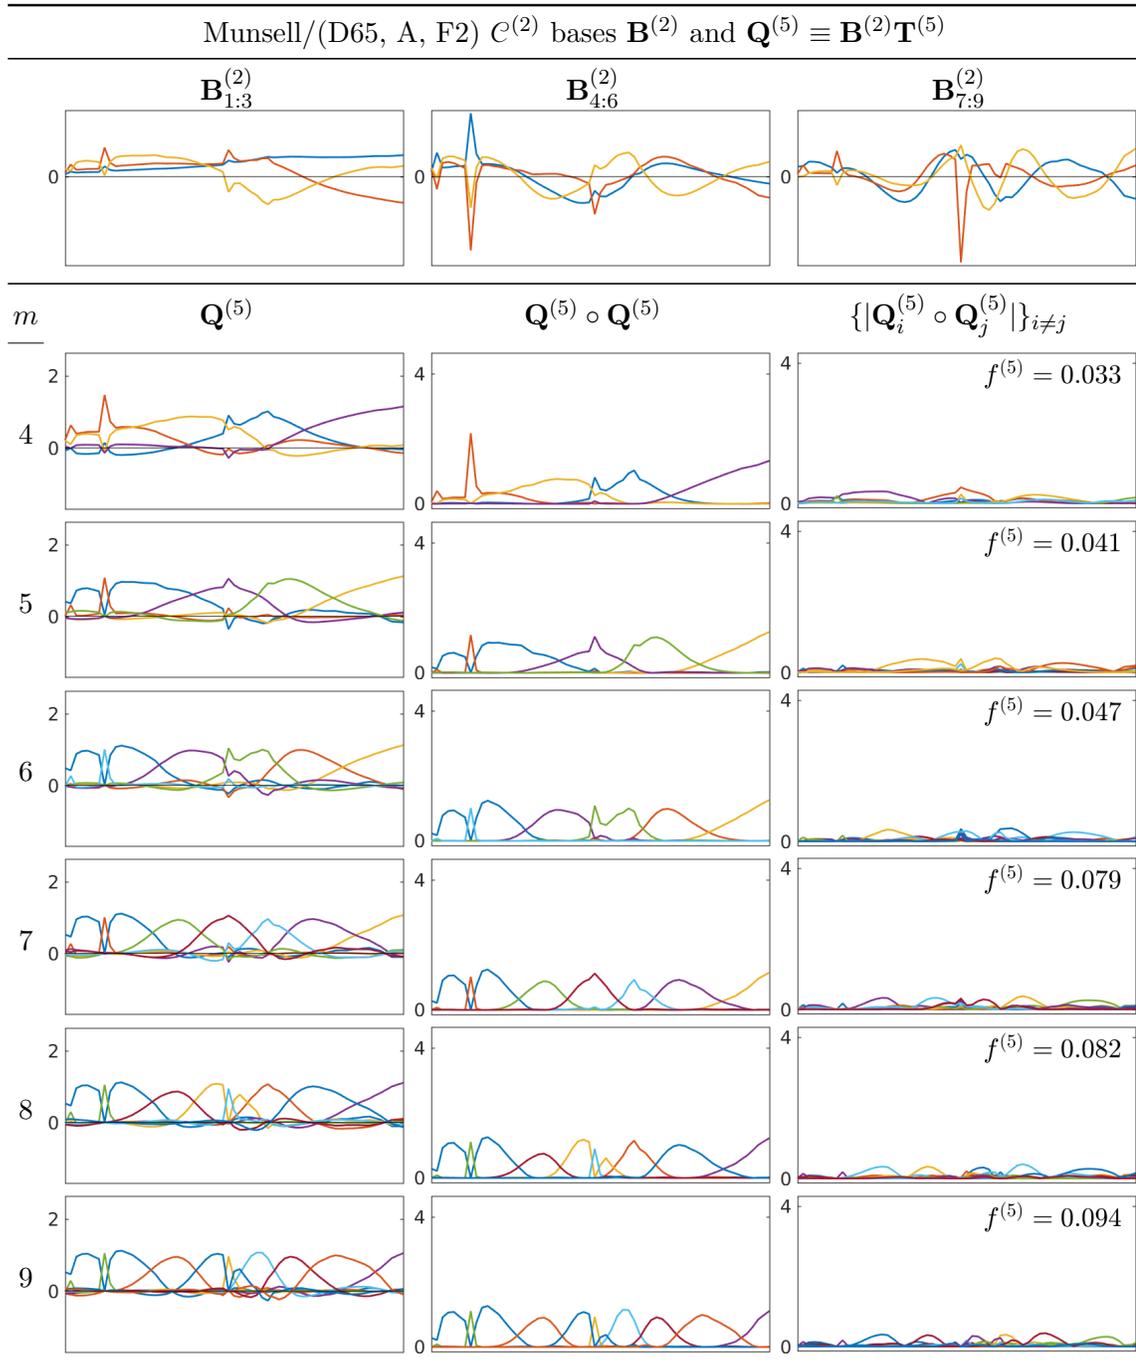


Figure 5.33: For Munsell $\mathcal{C}^{(2)}$ color signals formed with D65, A and F2, the orthonormal basis $\mathbf{B}^{(2)}$, the sharp bases $\mathbf{Q}^{(5)} \equiv \mathbf{B}^{(2)}\mathbf{T}^{(5)}$ (for dimensionalities $m \equiv 4, \dots, 9$), their squares and the absolute values of the pairwise products of their columns.

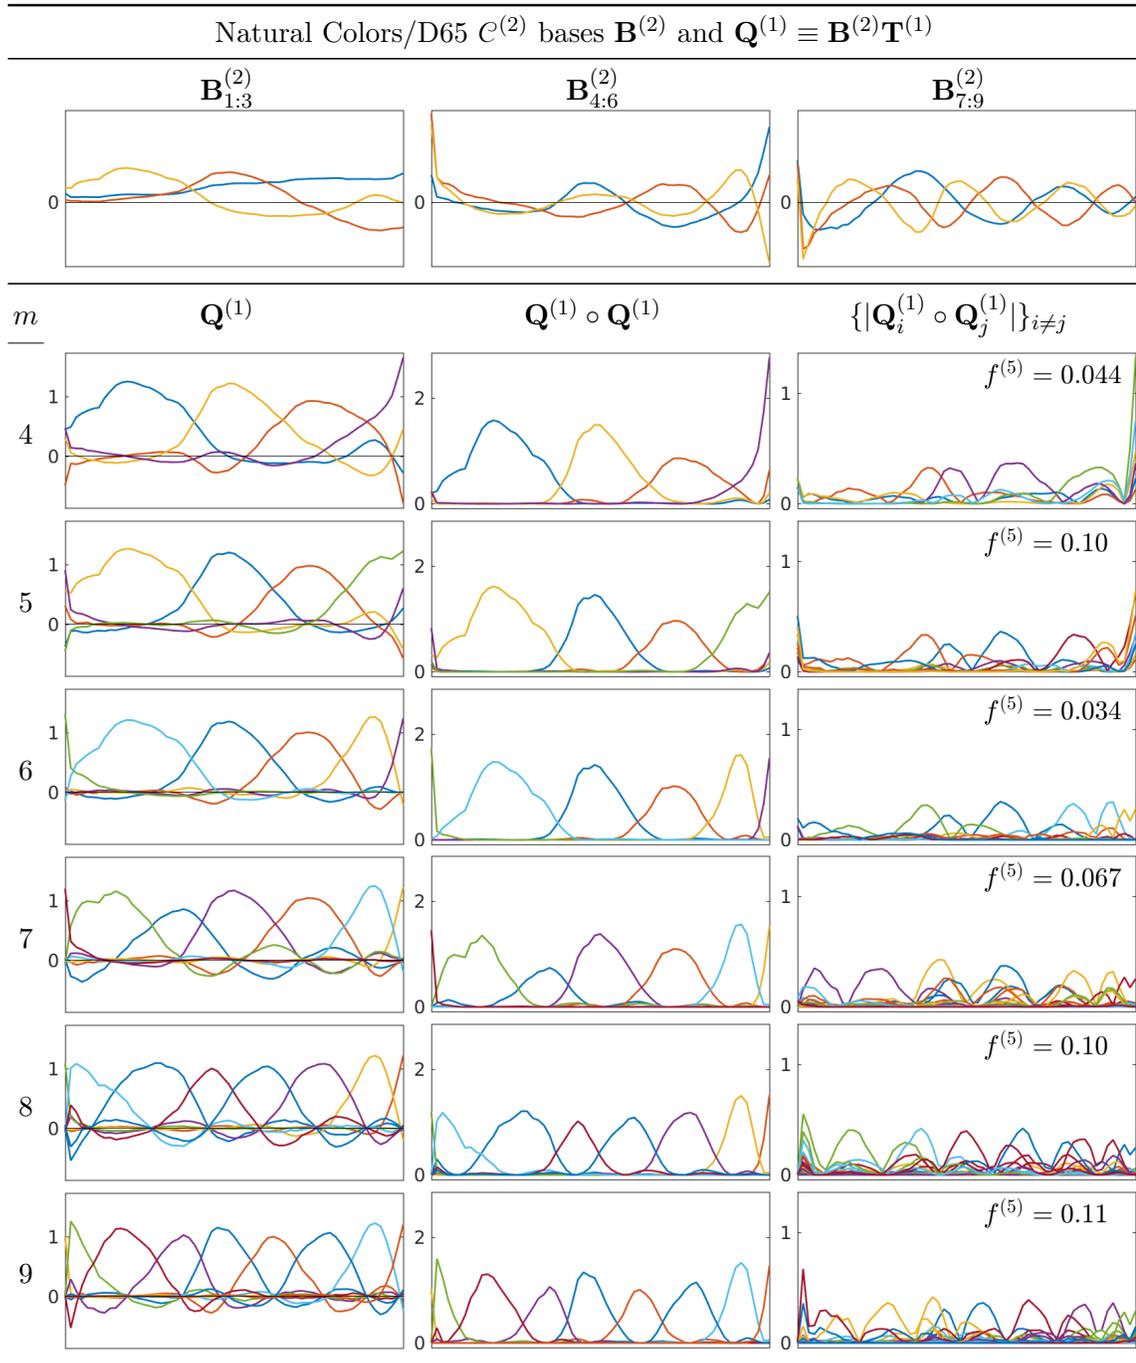


Figure 5.34: For Natural Colors $\mathcal{C}^{(2)}$ color signals formed with D65, the orthonormal basis $\mathbf{B}^{(2)}$, the sharp bases $\mathbf{Q}^{(1)} \equiv \mathbf{B}^{(2)}\mathbf{T}^{(1)}$ (for dimensionalities $m \equiv 4, \dots, 9$), their squares and the absolute values of the pairwise products of their columns.

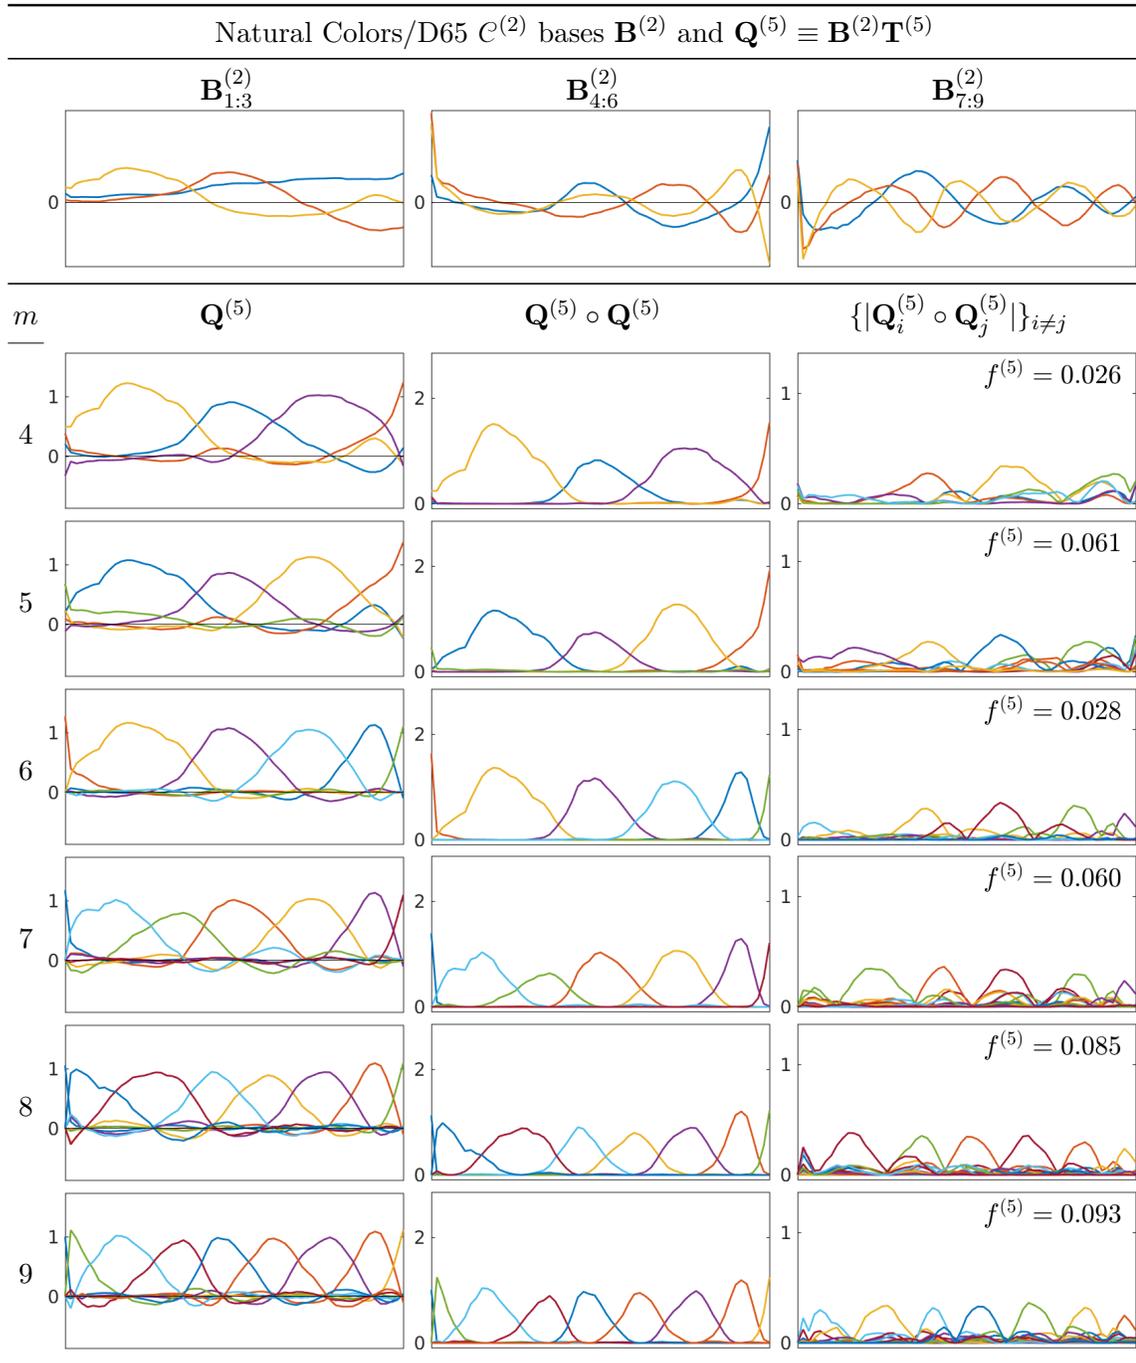


Figure 5.35: For Natural Colors $\mathcal{C}^{(2)}$ color signals formed with D65, the orthonormal basis $\mathbf{B}^{(2)}$, the sharp bases $\mathbf{Q}^{(5)} \equiv \mathbf{B}^{(2)}\mathbf{T}^{(5)}$ (for dimensionalities $m \equiv 4, \dots, 9$), their squares and the absolute values of the pairwise products of their columns.

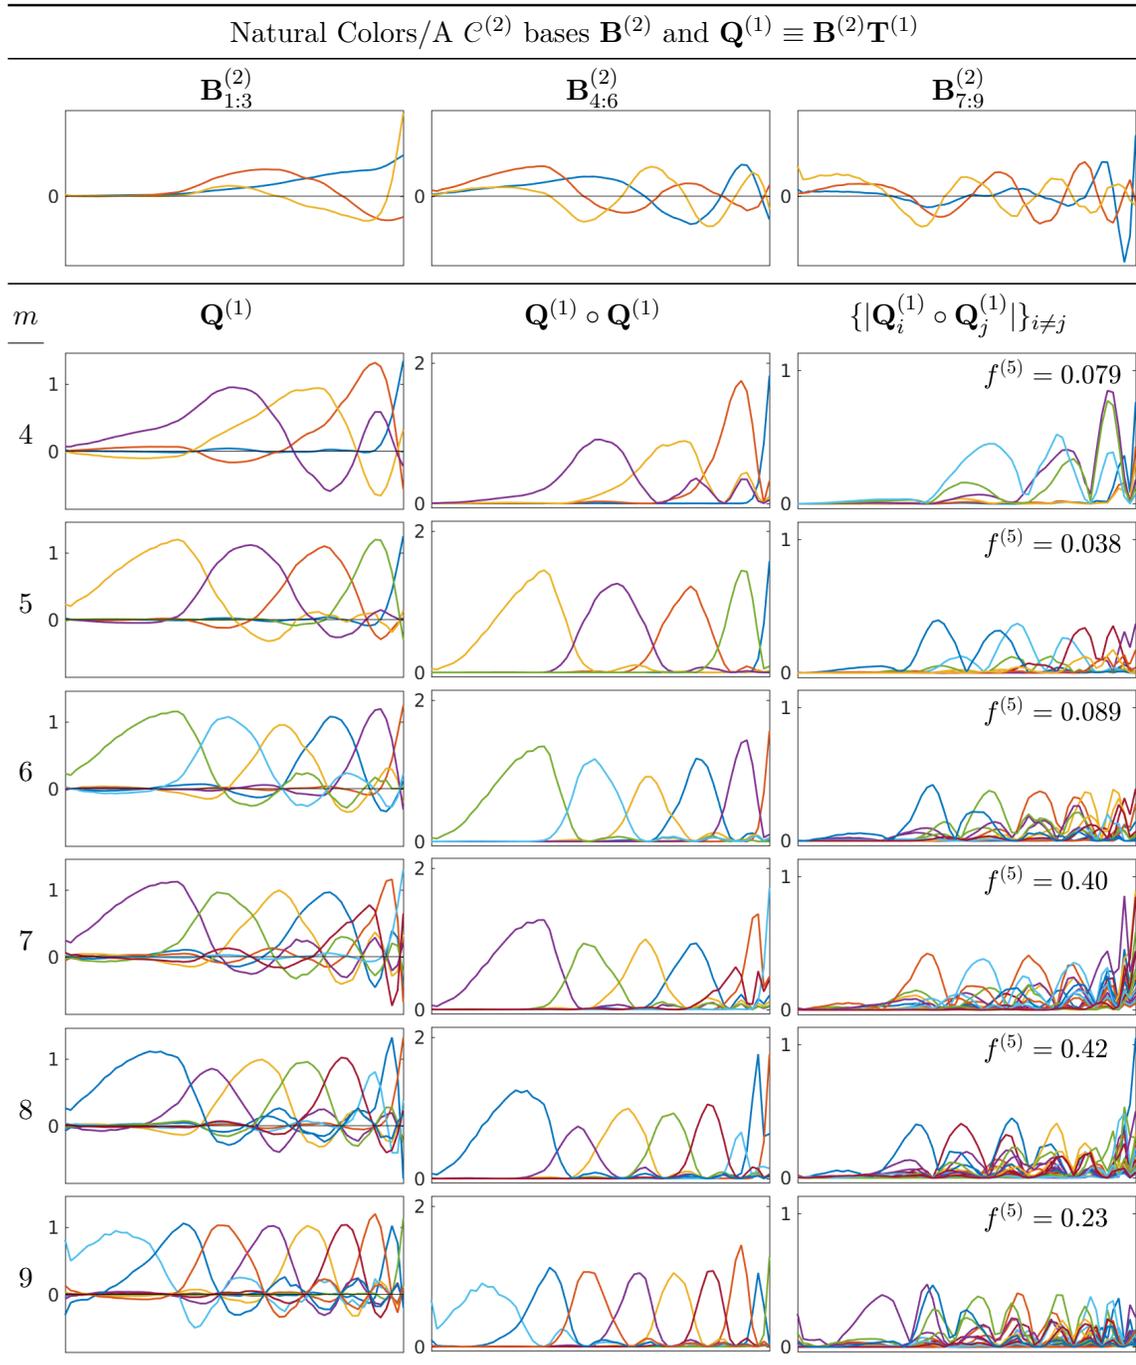


Figure 5.36: For Natural Colors $\mathcal{C}^{(2)}$ color signals formed with illuminant A, the orthonormal basis $\mathbf{B}^{(2)}$, the sharp bases $\mathbf{Q}^{(1)} \equiv \mathbf{B}^{(2)}\mathbf{T}^{(1)}$ (for dimensionalities $m \equiv 4, \dots, 9$), their squares and the absolute values of the pairwise products of their columns.

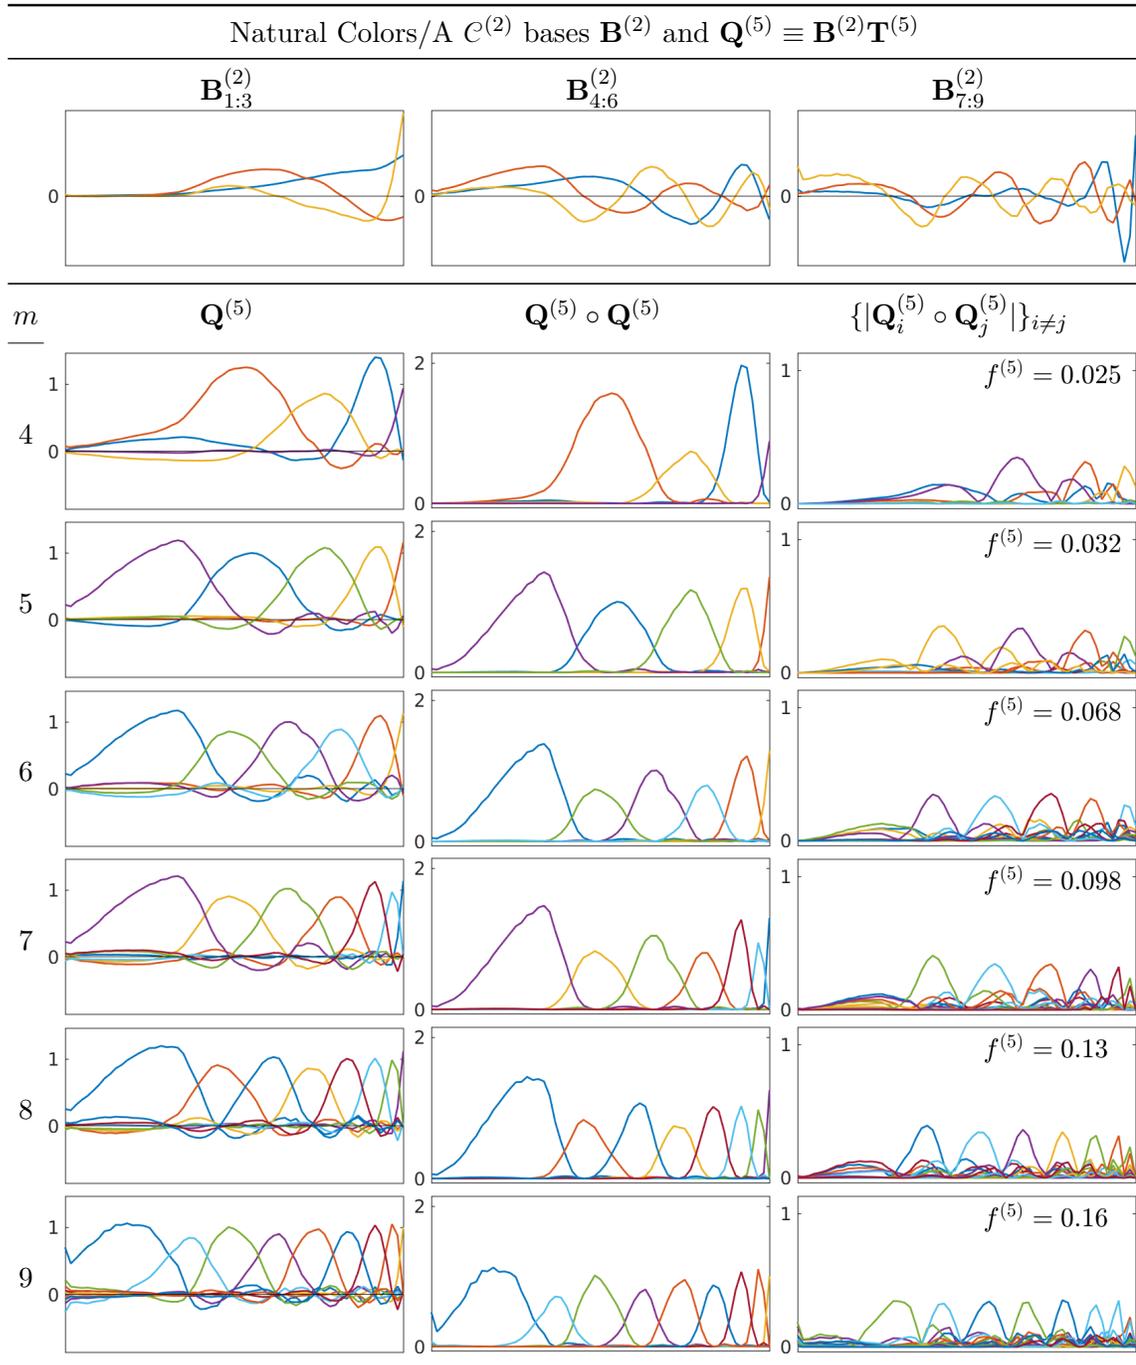


Figure 5.37: For Natural Colors $\mathcal{C}^{(2)}$ color signals formed with illuminant A, the orthonormal basis $\mathbf{B}^{(2)}$, the sharp bases $\mathbf{Q}^{(5)} \equiv \mathbf{B}^{(2)}\mathbf{T}^{(5)}$ (for dimensionalities $m \equiv 4, \dots, 9$), their squares and the absolute values of the pairwise products of their columns.

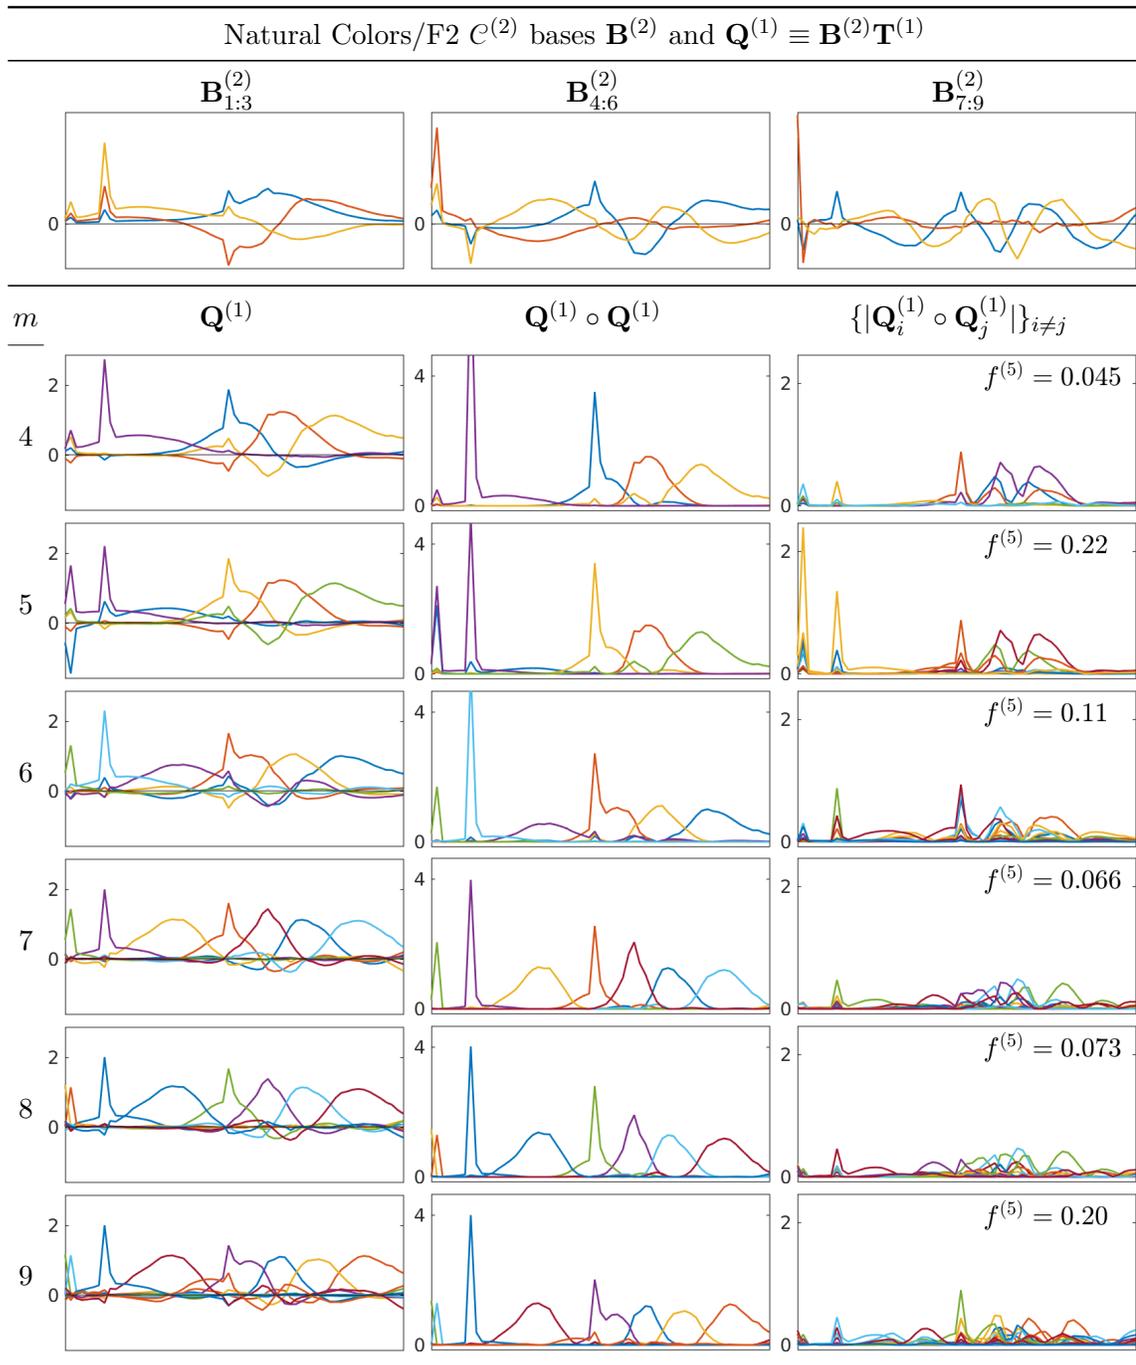


Figure 5.38: For Natural Colors $\mathcal{C}^{(2)}$ color signals formed with F2, the orthonormal basis $\mathbf{B}^{(2)}$, the sharp bases $\mathbf{Q}^{(1)} \equiv \mathbf{B}^{(2)}\mathbf{T}^{(1)}$ (for dimensionalities $m \equiv 4, \dots, 9$), their squares and the absolute values of the pairwise products of their columns.

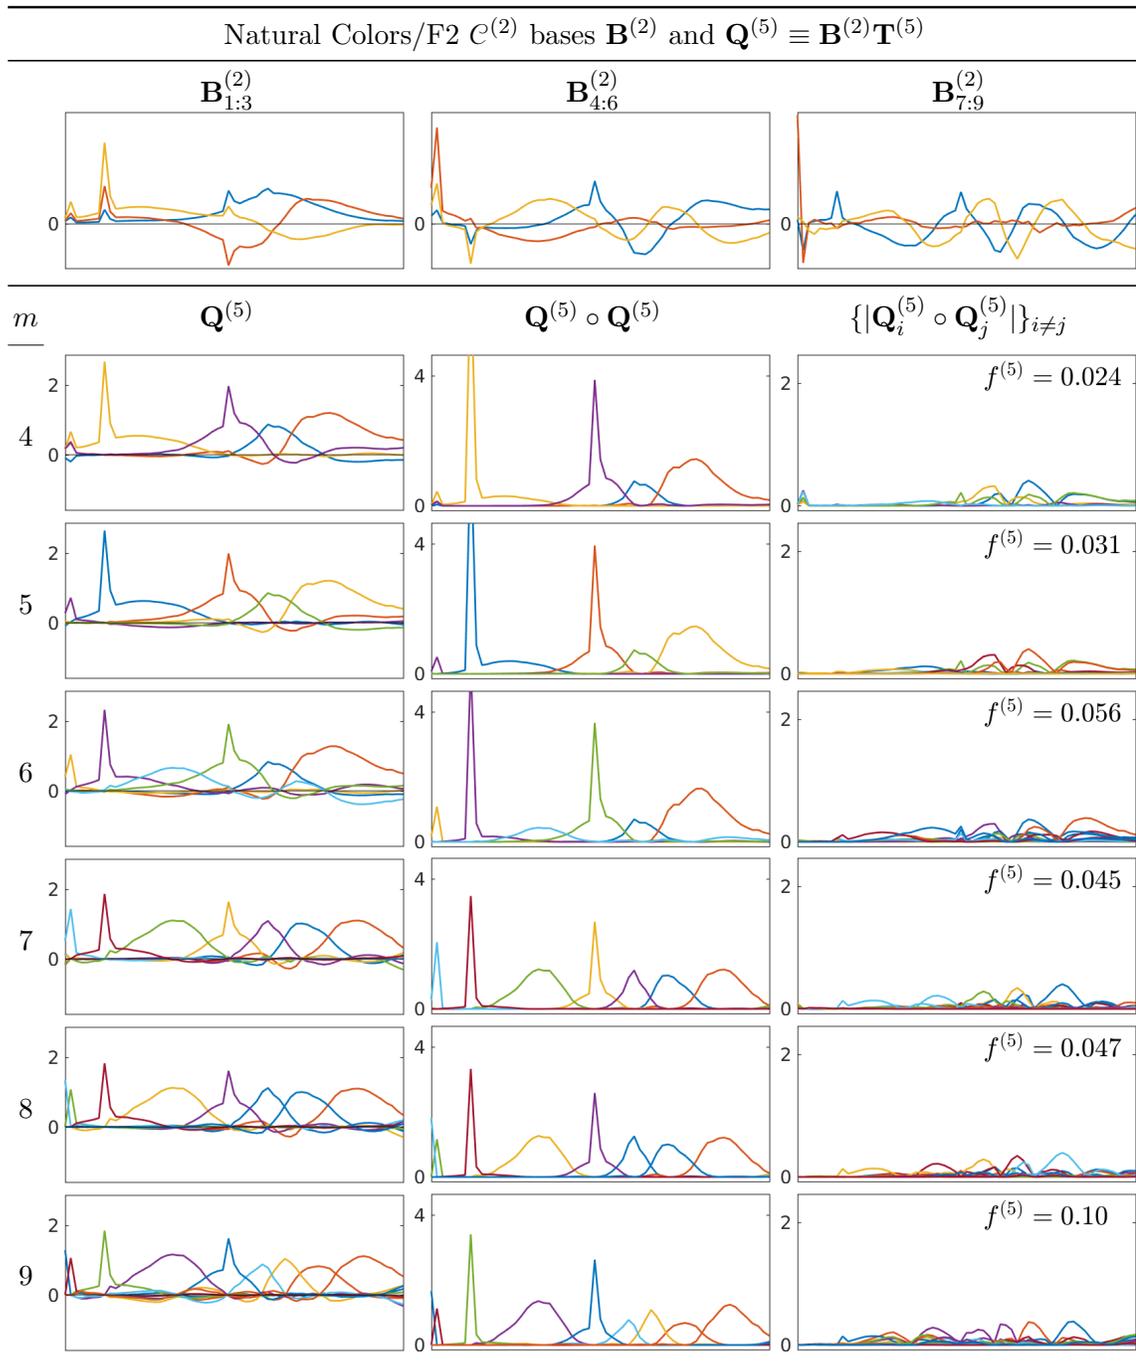


Figure 5.39: For Natural Colors $\mathcal{C}^{(2)}$ color signals formed with F2, the orthonormal basis $\mathbf{B}^{(2)}$, the sharp bases $\mathbf{Q}^{(5)} \equiv \mathbf{B}^{(2)}\mathbf{T}^{(5)}$ (for dimensionalities $m \equiv 4, \dots, 9$), their squares and the absolute values of the pairwise products of their columns.

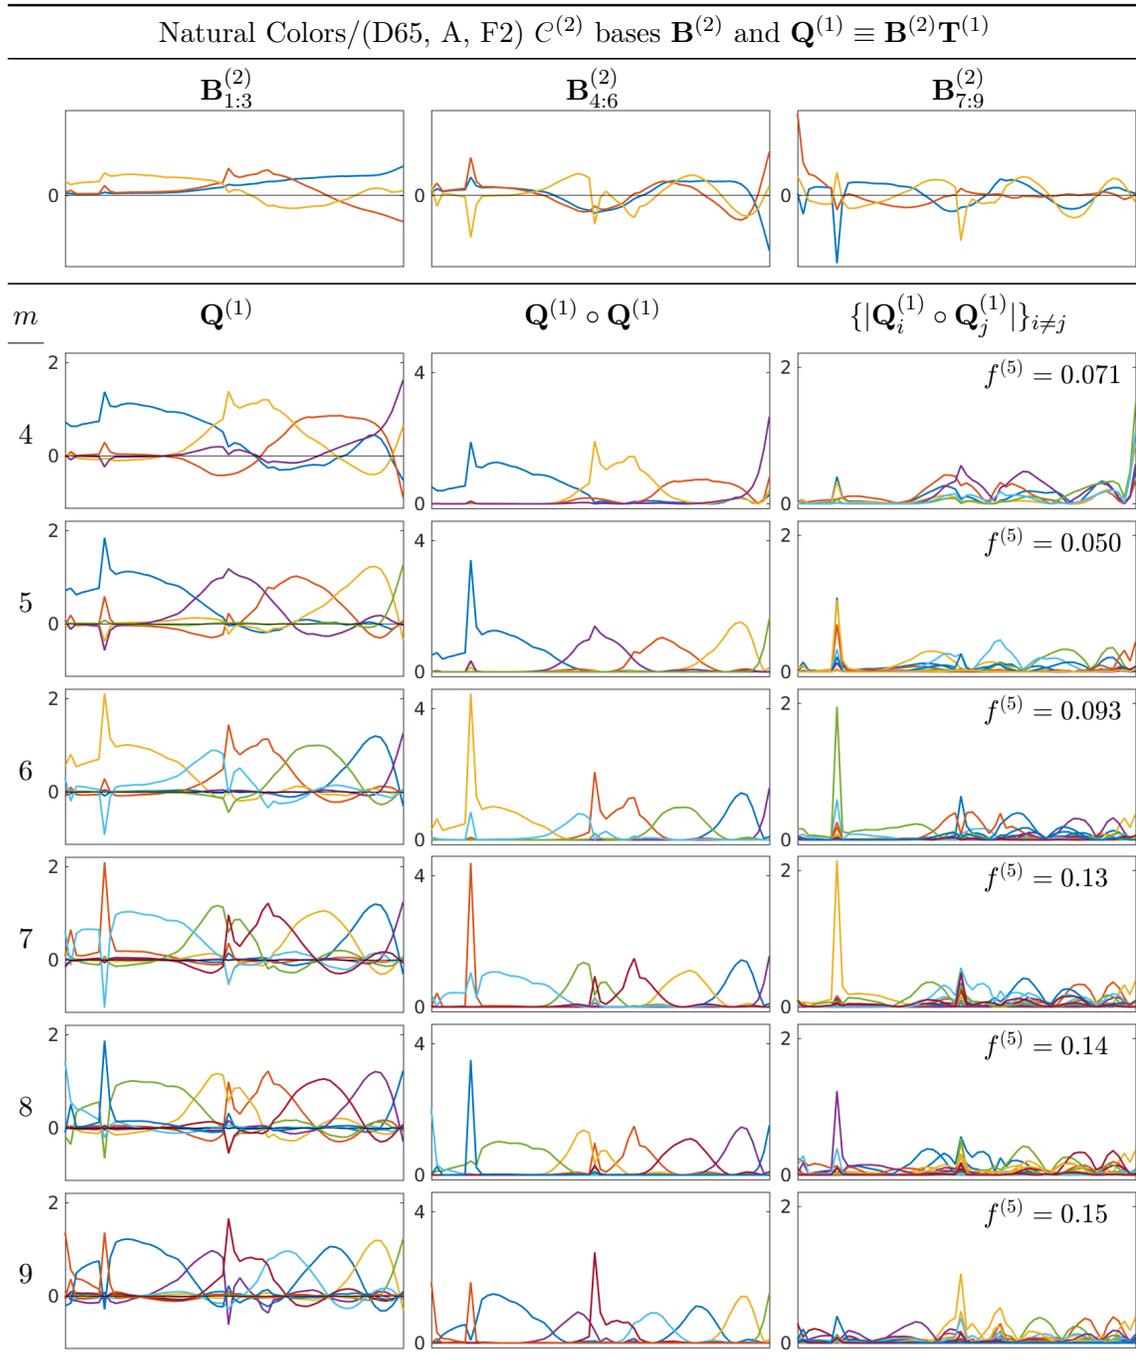


Figure 5.40: For Natural Colors $\mathcal{C}^{(2)}$ color signals formed with D65, A and F2, the orthonormal basis $\mathbf{B}^{(2)}$, the sharp bases $\mathbf{Q}^{(1)} \equiv \mathbf{B}^{(2)}\mathbf{T}^{(1)}$ (for dimensionalities $m \equiv 4, \dots, 9$), their squares and the absolute values of the pairwise products of their columns.

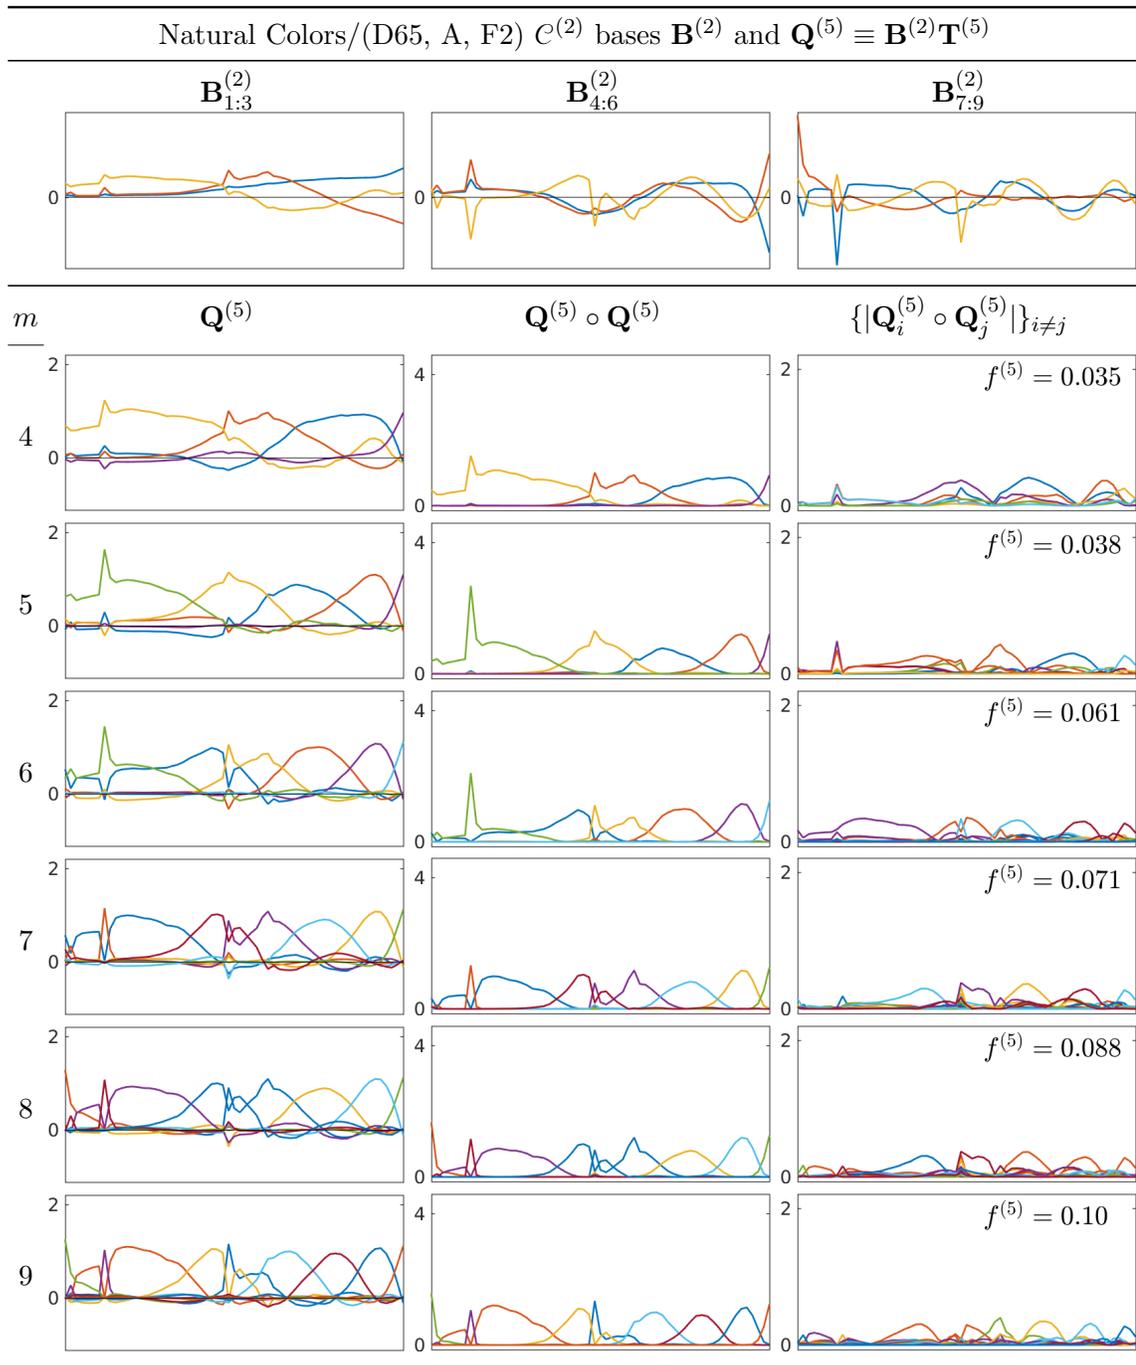


Figure 5.41: For Natural Colors $\mathcal{C}^{(2)}$ color signals formed with D65, A and F2, the orthonormal basis $\mathbf{B}^{(2)}$, the sharp bases $\mathbf{Q}^{(5)} \equiv \mathbf{B}^{(2)}\mathbf{T}^{(5)}$ (for dimensionalities $m \equiv 4, \dots, 9$), their squares and the absolute values of the pairwise products of their columns.

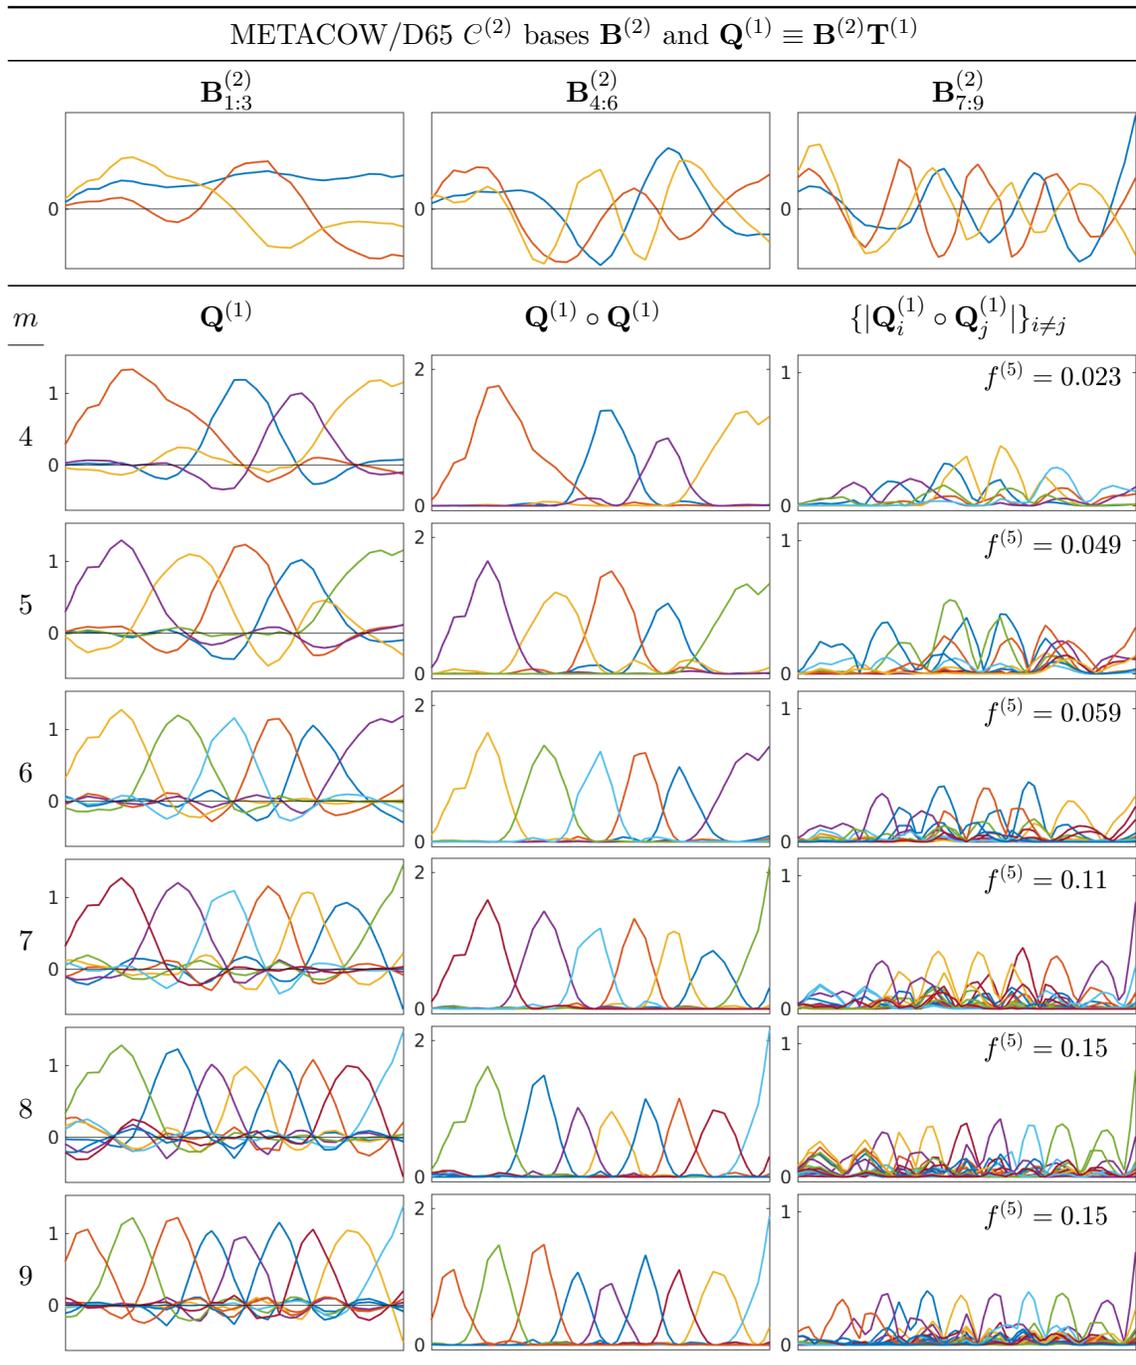


Figure 5.42: For METACOW $C^{(2)}$ color signals formed with D65, the orthonormal basis $\mathbf{B}^{(2)}$, the sharp bases $\mathbf{Q}^{(1)} \equiv \mathbf{B}^{(2)}\mathbf{T}^{(1)}$ (for dimensionalities $m \equiv 4, \dots, 9$), their squares and the absolute values of the pairwise products of their columns.

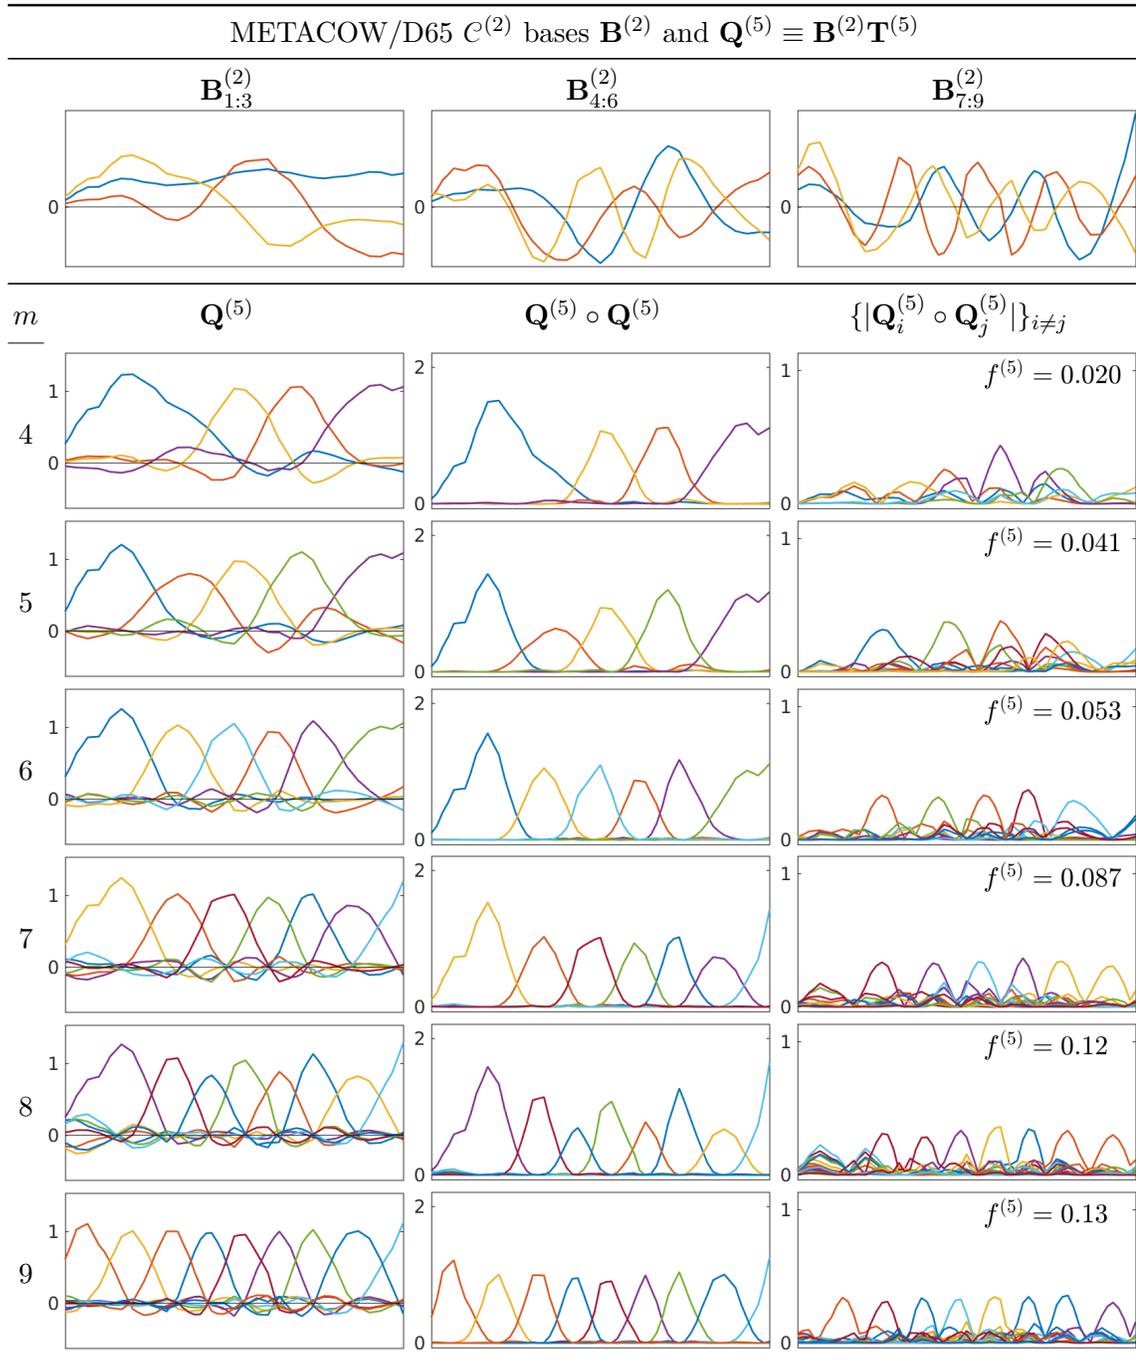


Figure 5.43: For METACOW $C^{(2)}$ color signals formed with D65, the orthonormal basis $\mathbf{B}^{(2)}$, the sharp bases $\mathbf{Q}^{(5)} \equiv \mathbf{B}^{(2)}\mathbf{T}^{(5)}$ (for dimensionalities $m \equiv 4, \dots, 9$), their squares and the absolute values of the pairwise products of their columns.

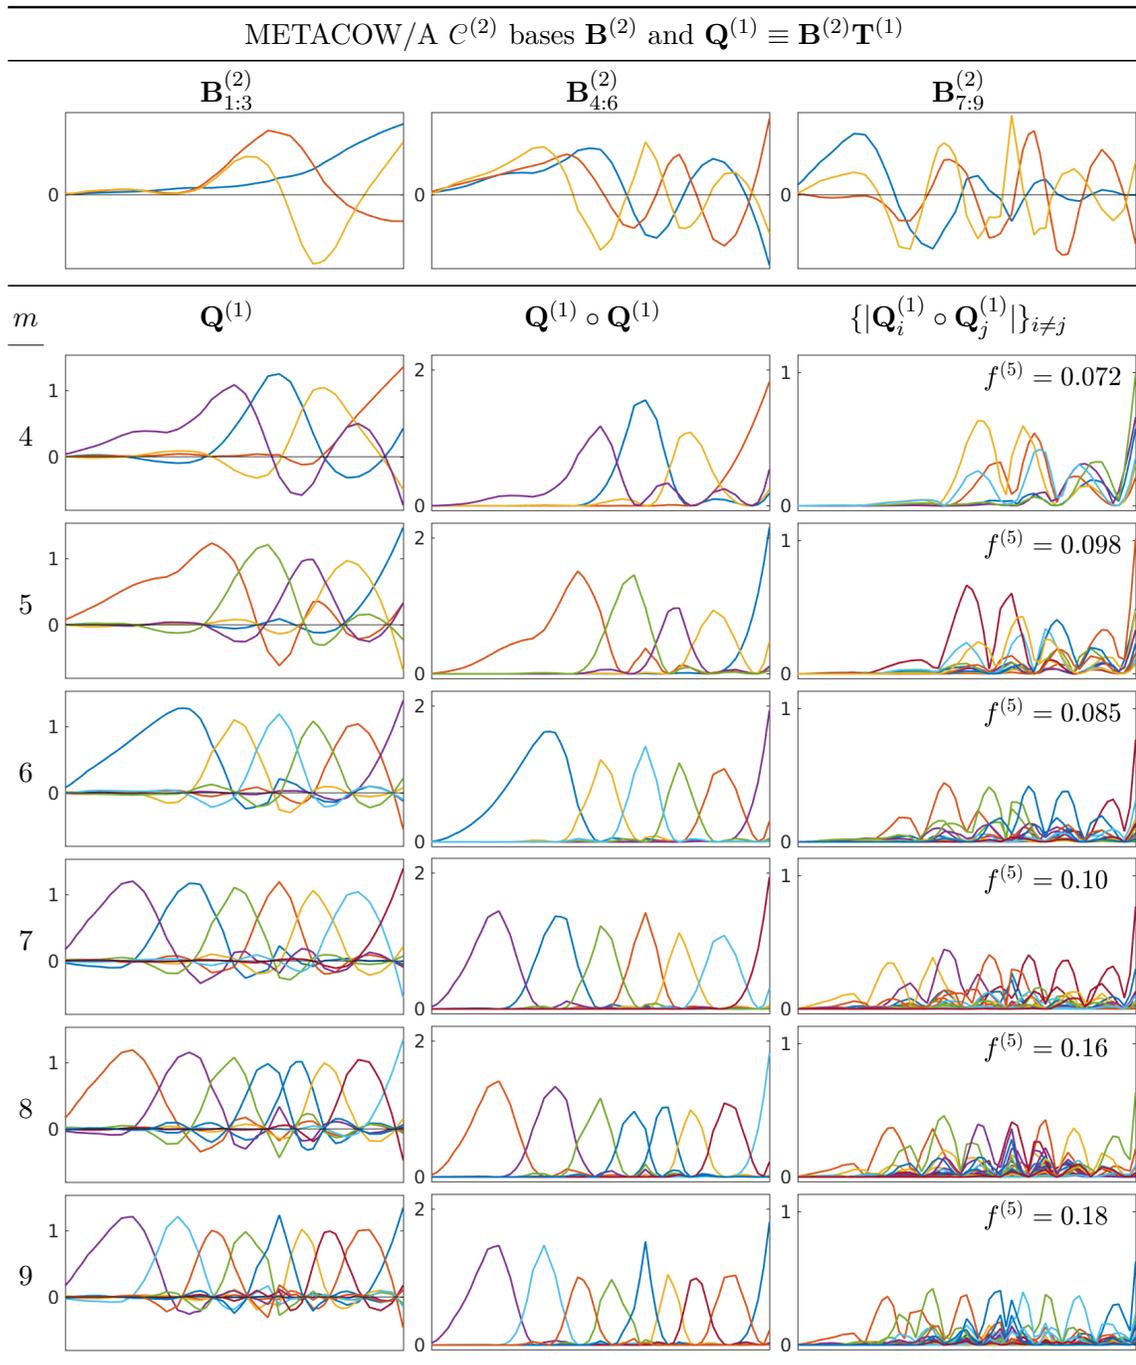


Figure 5.44: For METACOW $C^{(2)}$ color signals formed with illuminant A, the orthonormal basis $\mathbf{B}^{(2)}$, the sharp bases $\mathbf{Q}^{(1)} \equiv \mathbf{B}^{(2)}\mathbf{T}^{(1)}$ (for dimensionalities $m \equiv 4, \dots, 9$), their squares and the absolute values of the pairwise products of their columns.

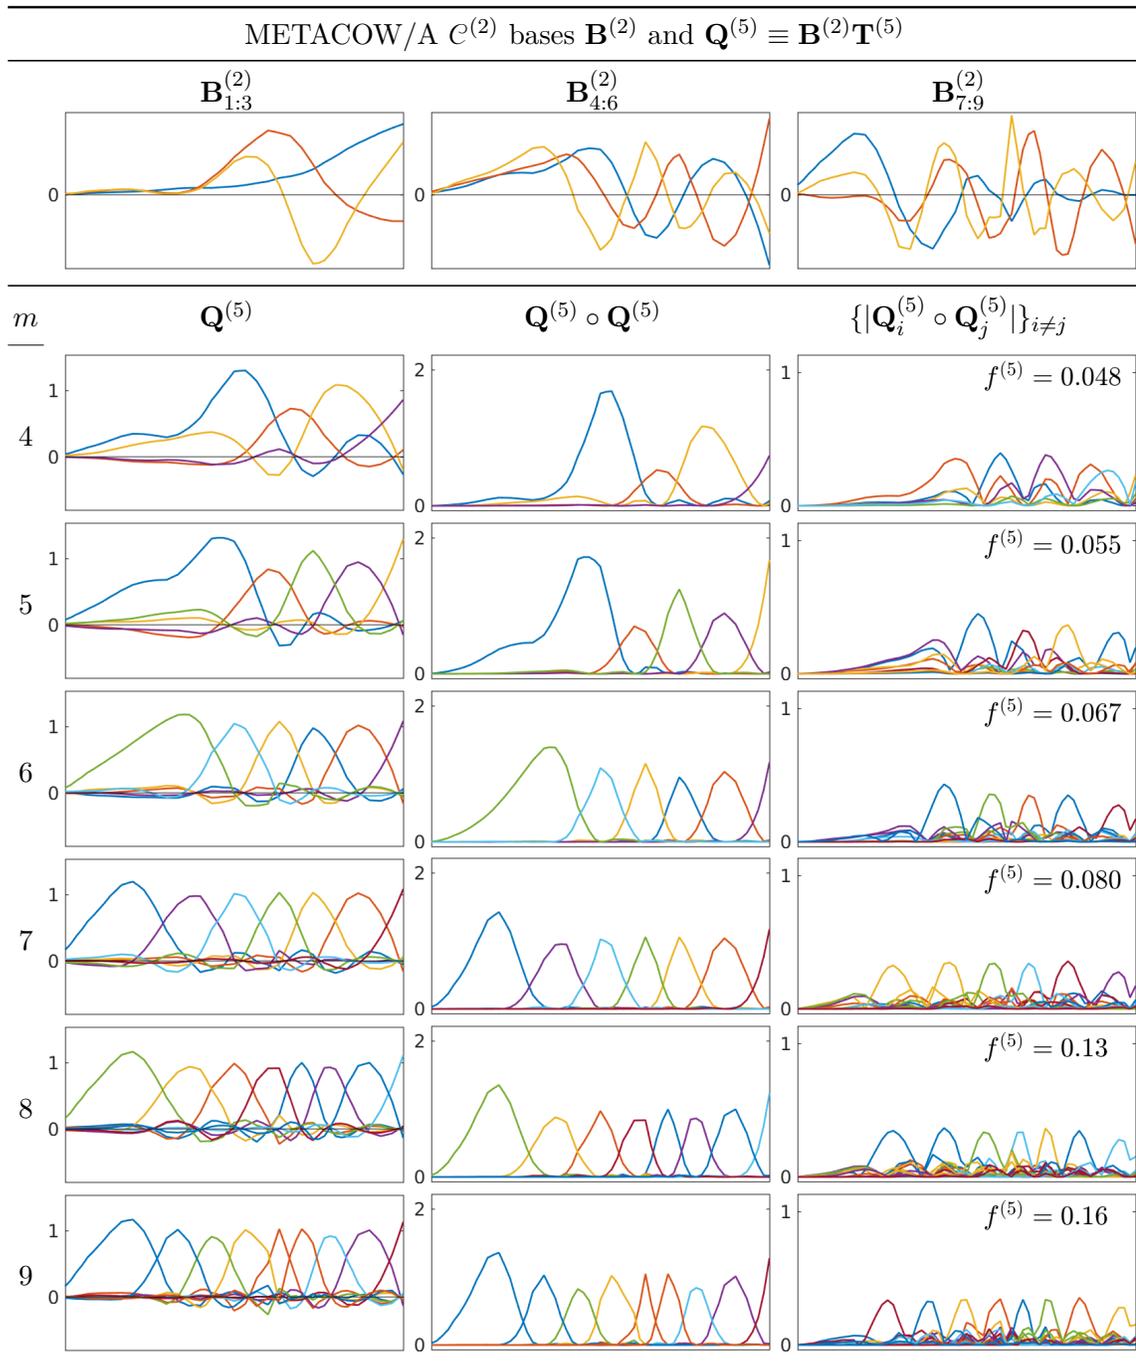


Figure 5.45: For METACOW $C^{(2)}$ color signals formed with illuminant A, the orthonormal basis $\mathbf{B}^{(2)}$, the sharp bases $\mathbf{Q}^{(5)} \equiv \mathbf{B}^{(2)}\mathbf{T}^{(5)}$ (for dimensionalities $m \equiv 4, \dots, 9$), their squares and the absolute values of the pairwise products of their columns.

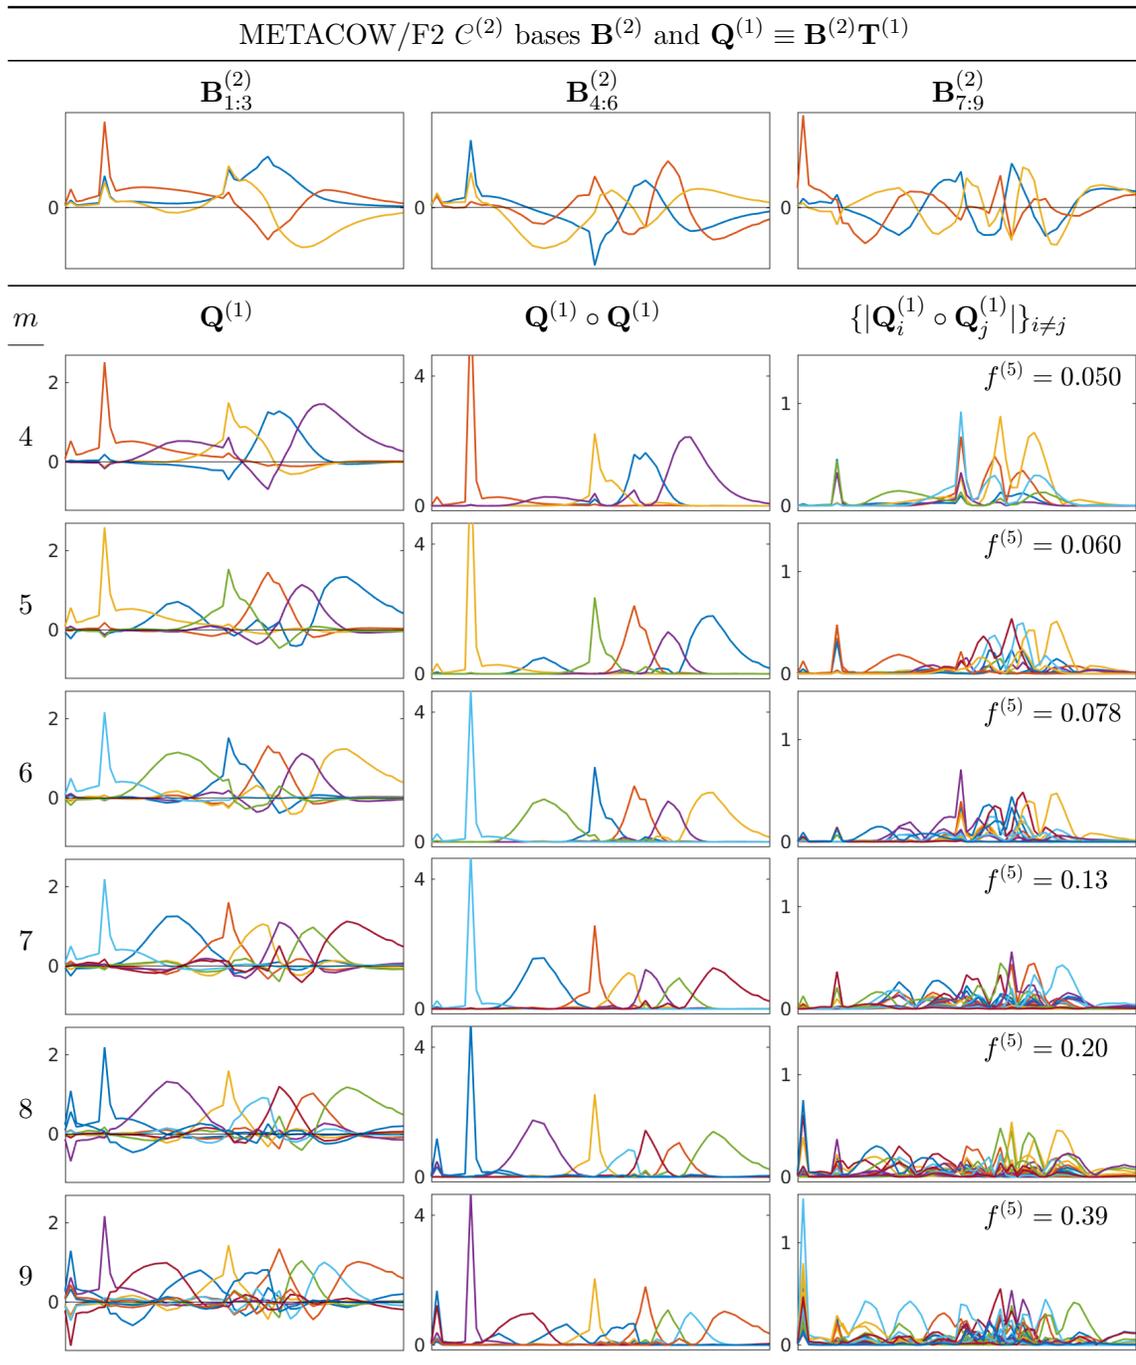


Figure 5.46: For METACOW $\mathcal{C}^{(2)}$ color signals formed with F2, the orthonormal basis $\mathbf{B}^{(2)}$, the sharp bases $\mathbf{Q}^{(1)} \equiv \mathbf{B}^{(2)}\mathbf{T}^{(1)}$ (for dimensionalities $m \equiv 4, \dots, 9$), their squares and the absolute values of the pairwise products of their columns.

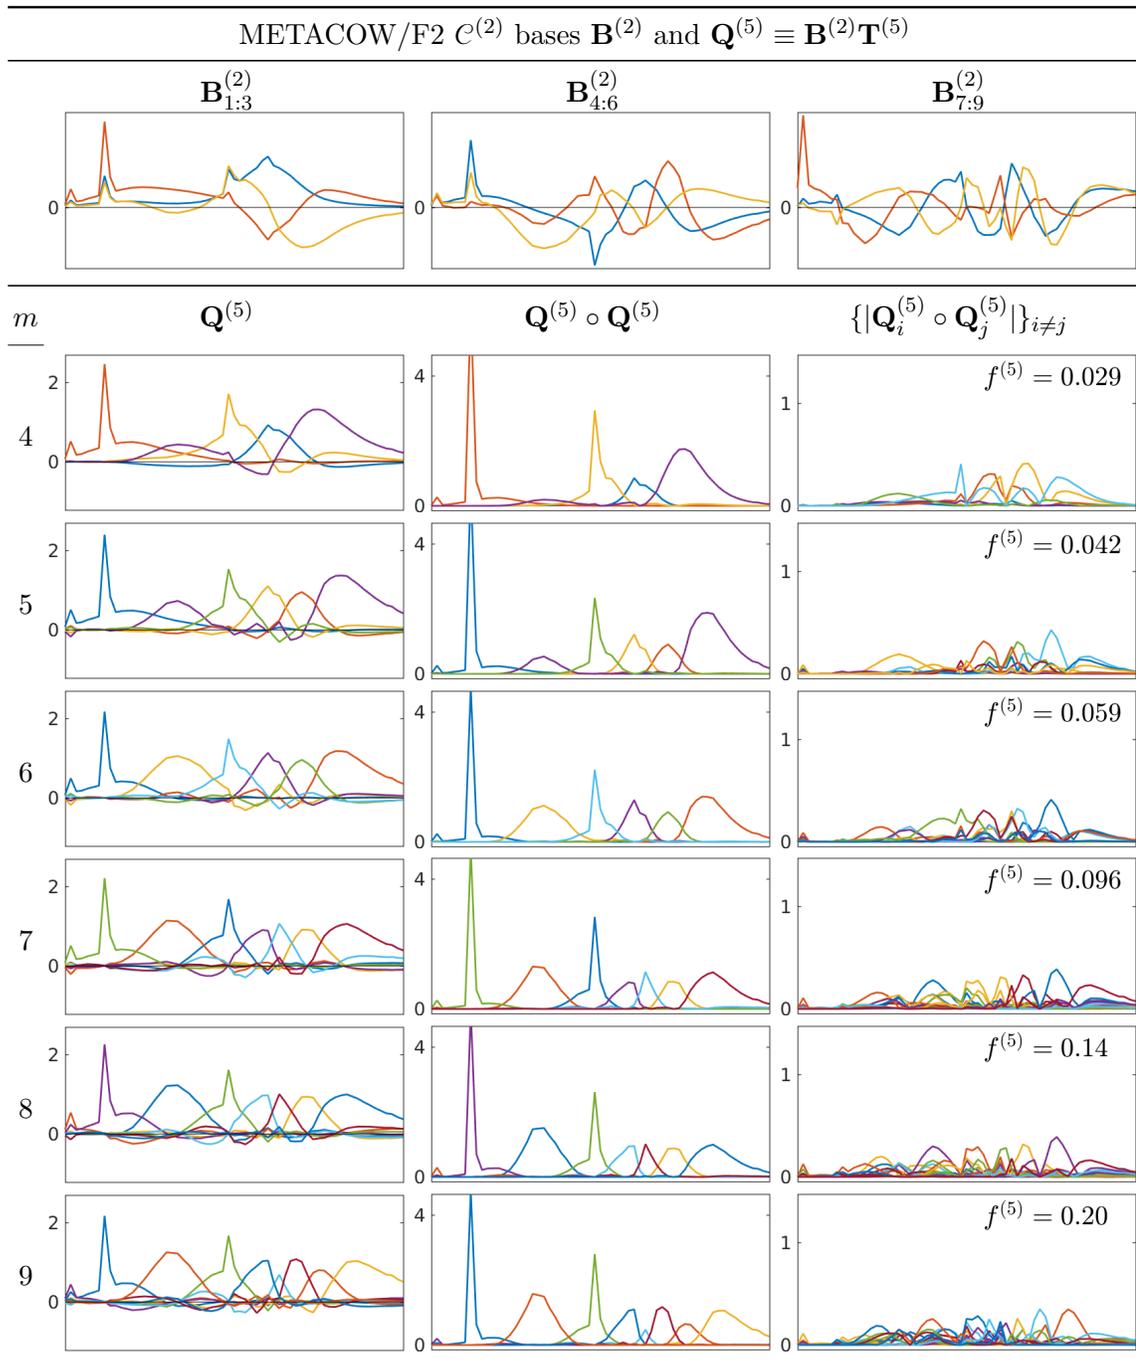


Figure 5.47: For METACOW $\mathcal{C}^{(2)}$ color signals formed with F2, the orthonormal basis $\mathbf{B}^{(2)}$, the sharp bases $\mathbf{Q}^{(5)} \equiv \mathbf{B}^{(2)}\mathbf{T}^{(5)}$ (for dimensionalities $m \equiv 4, \dots, 9$), their squares and the absolute values of the pairwise products of their columns.

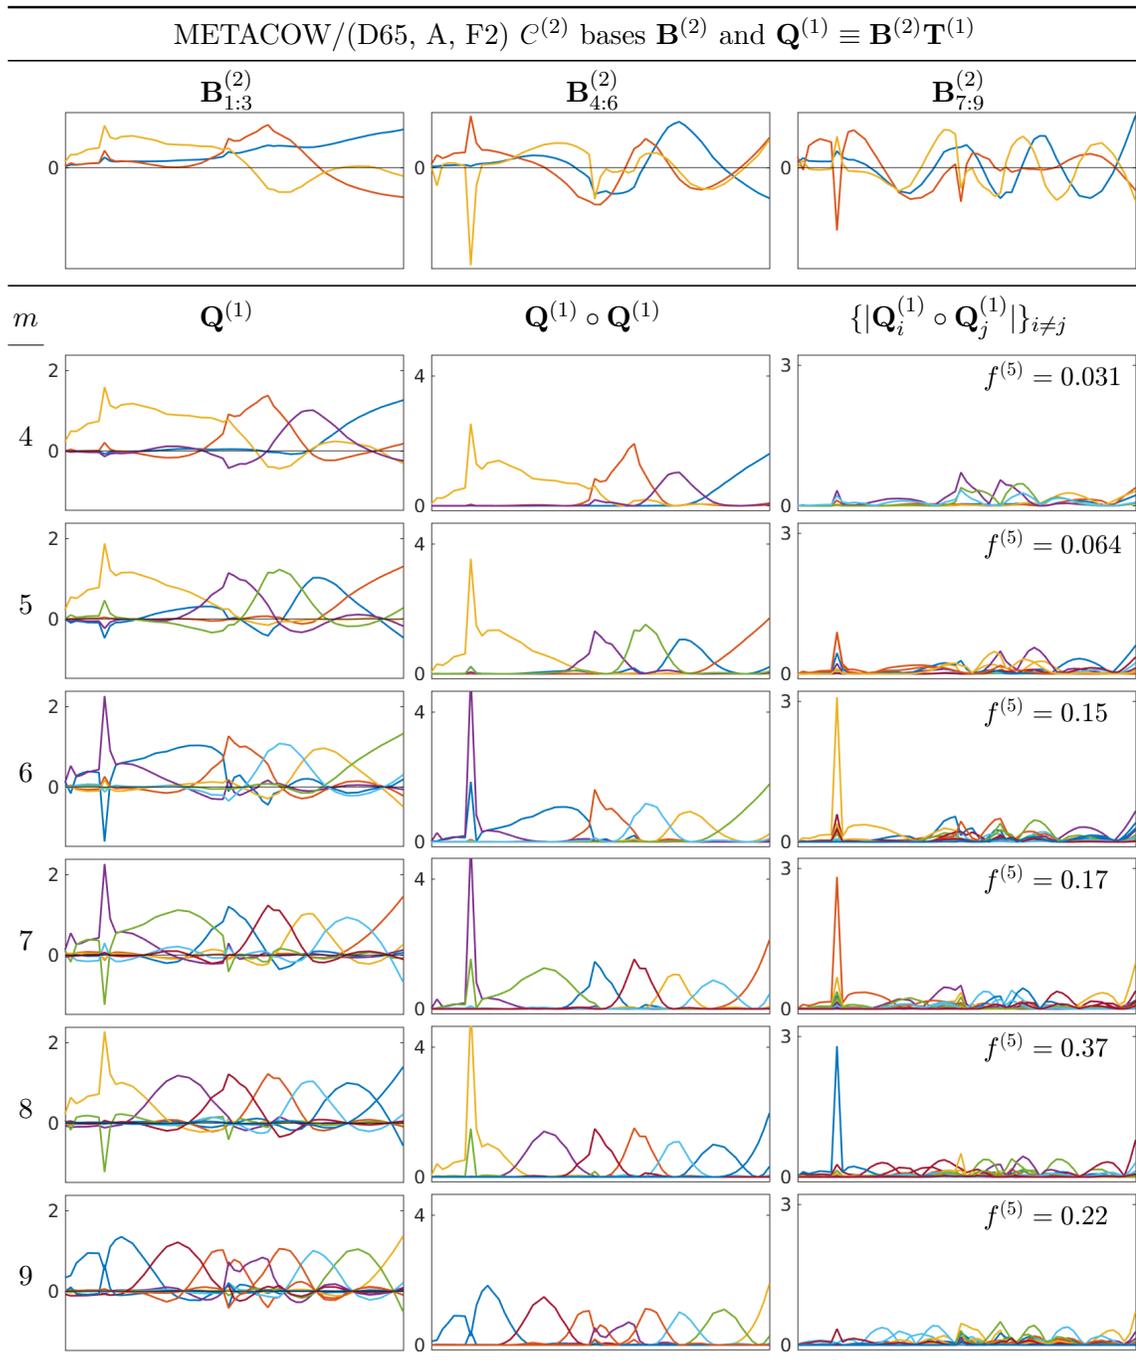


Figure 5.48: For METACOW $\mathcal{C}^{(2)}$ color signals formed with D65, A and F2, the orthonormal basis $\mathbf{B}^{(2)}$, the sharp bases $\mathbf{Q}^{(1)} \equiv \mathbf{B}^{(2)}\mathbf{T}^{(1)}$ (for dimensionalities $m \equiv 4, \dots, 9$), their squares and the absolute values of the pairwise products of their columns.

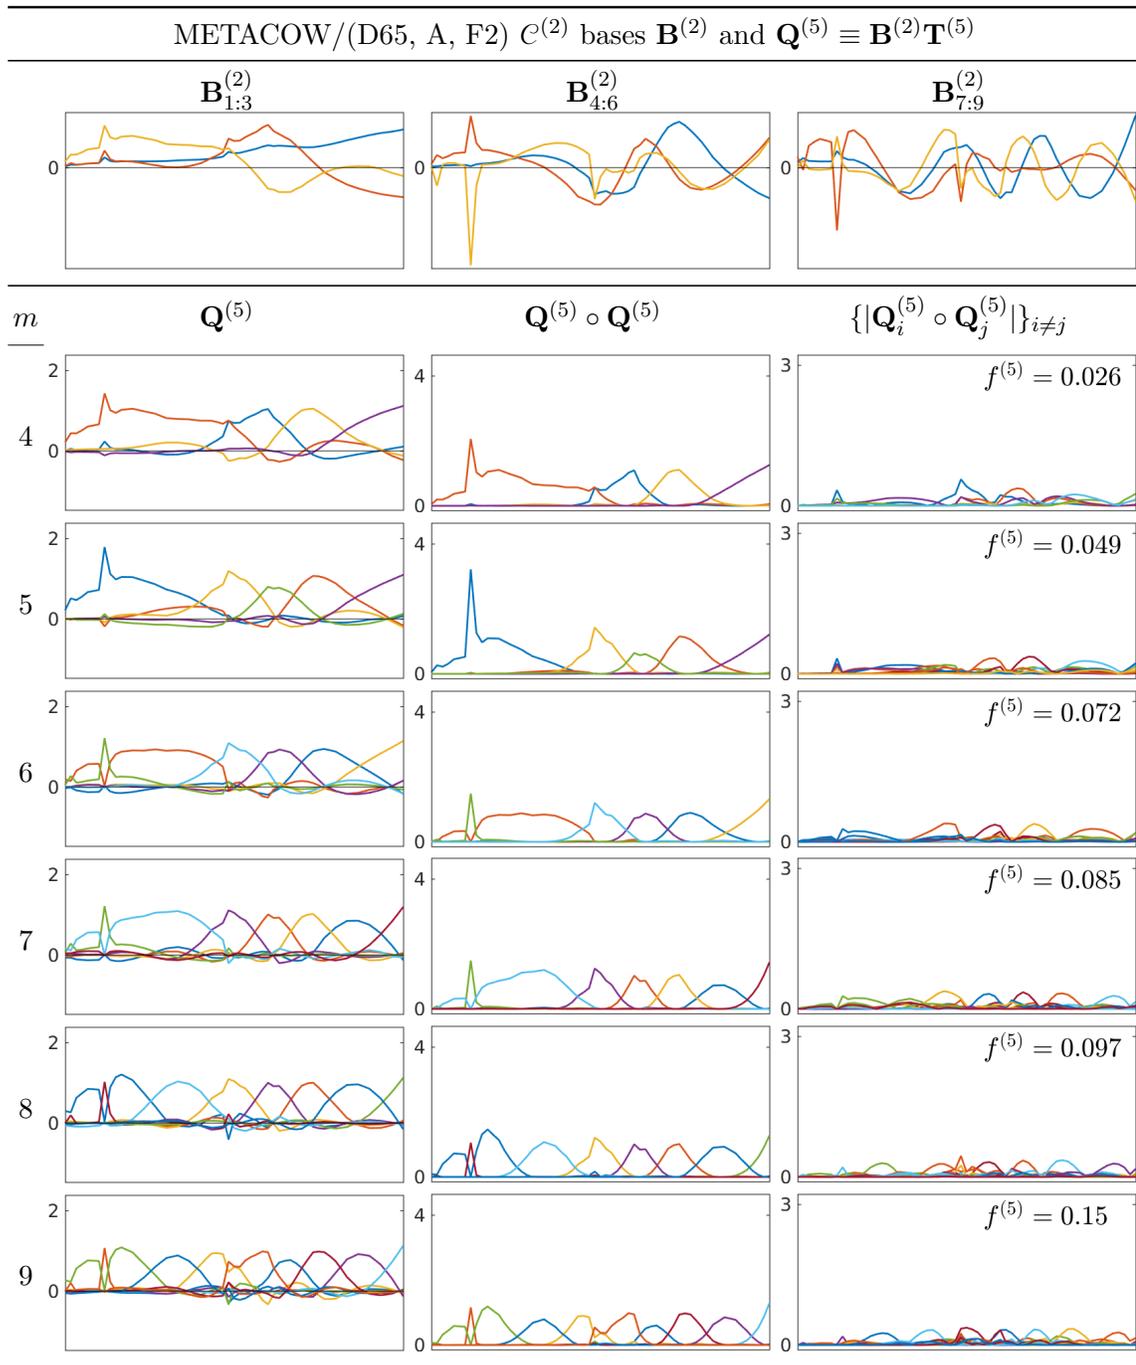


Figure 5.49: For METACOW $\mathcal{C}^{(2)}$ color signals formed with D65, A and F2, the orthonormal basis $\mathbf{B}^{(2)}$, the sharp bases $\mathbf{Q}^{(5)} \equiv \mathbf{B}^{(2)}\mathbf{T}^{(5)}$ (for dimensionalities $m \equiv 4, \dots, 9$), their squares and the absolute values of the pairwise products of their columns.

Chapter 6

Results

Accuracy and efficiency are competing goals of almost any approximation method. In Section 6.2 real-time rendering performance is used as a measure of the latter, which, for a given choice of the dimensionality m of a basis \mathbf{B} , is the same for all methods of finding \mathbf{T} described in the preceding chapter. In the next section the accuracy of these methods is reported, measured as spectral errors and the resulting color differences.

6.1 Accuracy

There are various ways to measure the accuracy of approximated spectra, several of which are considered and compared by Imai et al. [80] and Viggiano [167]. Because spectra of smaller magnitude make correspondingly smaller contributions to a rendered scene, our methods seek to minimize the absolute rather than the relative approximation error. For this reason, we measure the accuracy of an SPD $\mathbf{C} \equiv \mathbf{S}^{(k)} \circ \dots \circ \mathbf{S}^{(1)} \circ \mathbf{E}$ approximated by its projection $\hat{\mathbf{C}} \equiv \mathbf{B}\mathbf{B}^T\mathbf{C}$ onto \mathbf{B} or by multiplying coefficients of a sharp basis $\mathbf{Q} \equiv \mathbf{B}\mathbf{T}$,

$$\tilde{\mathbf{C}} \equiv \mathbf{Q}(\tilde{\mathbf{s}}^{(k)} \circ \dots \circ \tilde{\mathbf{s}}^{(1)} \circ \tilde{\mathbf{e}}),$$

with a normalized root-mean-square error (NRMSE) metric, intuitively scaled as a percentage of the color signal's largest possible RMS value, which is that of the light

source’s SPD:

$$\text{NRMSE}(\mathbf{C}, \mathbf{C}') \equiv 100 \cdot \frac{\|\mathbf{C} - \mathbf{C}'\|_2}{\|\mathbf{E}\|_2}.$$

Any negative components of the approximation \mathbf{C}' are clamped to zero to enforce physical plausibility. Since $\mathbf{B}^{(k)}$ is the best fit basis for a given set $\mathcal{C}^{(k)}$ and a given dimensionality m , the lowest possible sum of these errors is the residual error incurred by projection onto $\mathbf{B}^{(k)}$. That is,

$$\sum_{\mathbf{C} \in \mathcal{C}^{(k)}} \text{NRMSE}(\mathbf{C}, \tilde{\mathbf{C}}) \geq \sum_{\mathbf{C} \in \mathcal{C}^{(k)}} \text{NRMSE}(\mathbf{C}, \hat{\mathbf{C}}).$$

For SPDs generated by more than k reflections, this inequality will hold approximately if the subspace spanned by $\mathbf{B}^{(k)}$ remains a good fit, which, for typical reflectance spectra and $k \equiv 2$, should be expected.

We measure the effect of spectral error on color accuracy with the CIEDE2000 color difference formula [33], denoted ΔE_{00} , which, like the NRMSE error metric, is applied after clamping negative components of approximated spectra to zero. A commonly cited rule of thumb is that when comparing solid color patches a difference of 1 is just noticeable, but differences of 2.3 ± 1.3 have been measured as imperceptible [113]. When colors are perceived in the context of an image instead of as patches, even larger differences may be imperceptible [27]. In any case, the color differences reported here should not be interpreted as absolute measures of color accuracy, because the viewing conditions for which they are defined are quite different. In particular, the CIEDE2000 formula measures color differences of large patches on a uniform gray background with D65 illumination. Nevertheless, when applied to pixel color differences, this formula provides a useful estimation of the perceptibility of differences between images, and in at least one study it was found that differences between images with an average pixel color difference of 2.15 were not noticeable [156]. Understood as heuristic in this context, color differences provide an approximate, relative scaling of the effects of spectral errors on the perception of color images.

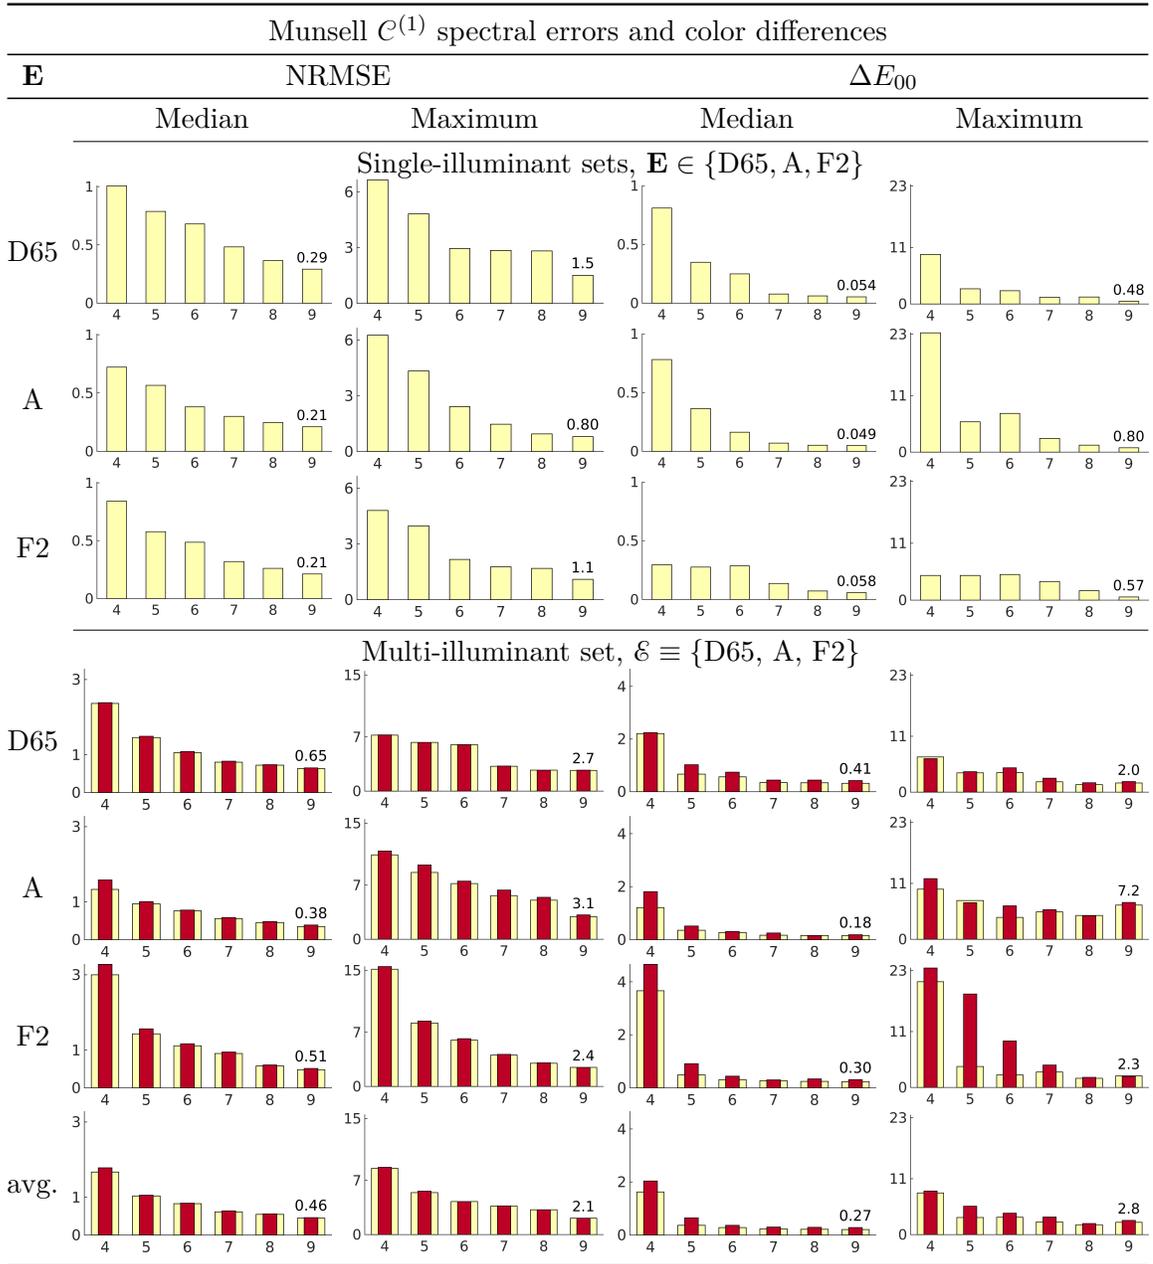


Figure 6.1: For dimensionalities $m \equiv 4, \dots, 9$, the spectral errors and color differences incurred when computing Munsell $\mathcal{C}^{(1)}$ reflections of D65, A, F2, or all three. In the latter case, errors incurred by reflecting the average all three illuminants are shown in the bottom row. For single-illuminant sets, reflections are computed by multiplying coefficients of sharp basis $\mathbf{Q}^{(\mathbf{V})} \equiv \mathbf{B}^{(1)}\mathbf{T}^{(\mathbf{V})}$, where change of basis matrices $\mathbf{T}^{(\mathbf{V})}$ were found by diagonalizing $\mathbf{B}^{(1)\text{T}} \text{diag}(\mathbf{E})^{-1} \mathbf{B}^{(1)}$, as described in Section 5.1. The inner (darker) bars of the multi-illuminant set represent errors incurred using sharp basis $\mathbf{Q}^{(1)} \equiv \mathbf{B}^{(1)}\mathbf{T}^{(1)}$, where $\mathbf{T}^{(1)}$ minimizes $f^{(1)}$, the total approximation error, as described in Section 5.5. Outer (lighter) bars represent errors incurred by projection onto $\mathbf{B}^{(1)}$.

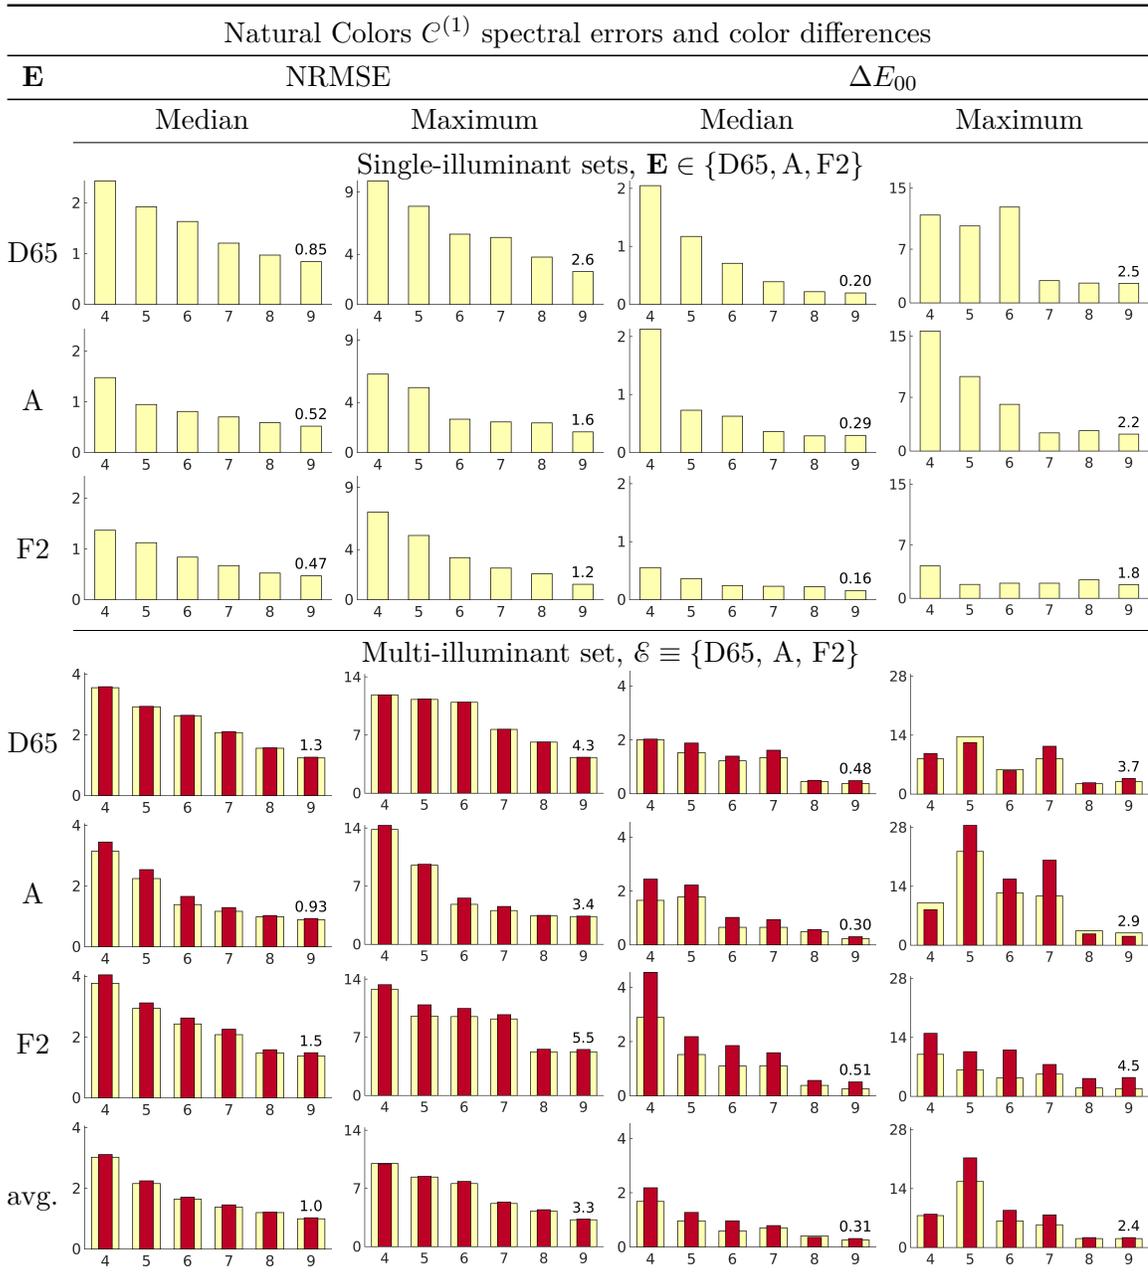


Figure 6.2: For dimensionalities $m \equiv 4, \dots, 9$, the spectral errors and color differences incurred when computing Natural Colors $\mathcal{C}^{(1)}$ reflections of D65, A, F2, or all three. In the latter case, errors incurred by reflecting the average all three illuminants are shown in the bottom row. For single-illuminant sets, reflections are computed by multiplying coefficients of sharp basis $\mathbf{Q}^{(\mathbf{V})} \equiv \mathbf{B}^{(1)}\mathbf{T}^{(\mathbf{V})}$, where change of basis matrices $\mathbf{T}^{(\mathbf{V})}$ were found by diagonalizing $\mathbf{B}^{(1)\text{T}} \text{diag}(\mathbf{E})^{-1}\mathbf{B}^{(1)}$, as described in Section 5.1. The inner (darker) bars of the multi-illuminant set represent errors incurred using sharp basis $\mathbf{Q}^{(1)} \equiv \mathbf{B}^{(1)}\mathbf{T}^{(1)}$, where $\mathbf{T}^{(1)}$ minimizes $f^{(1)}$, the total approximation error, as described in Section 5.5. Outer (lighter) bars represent errors incurred by projection onto $\mathbf{B}^{(1)}$.

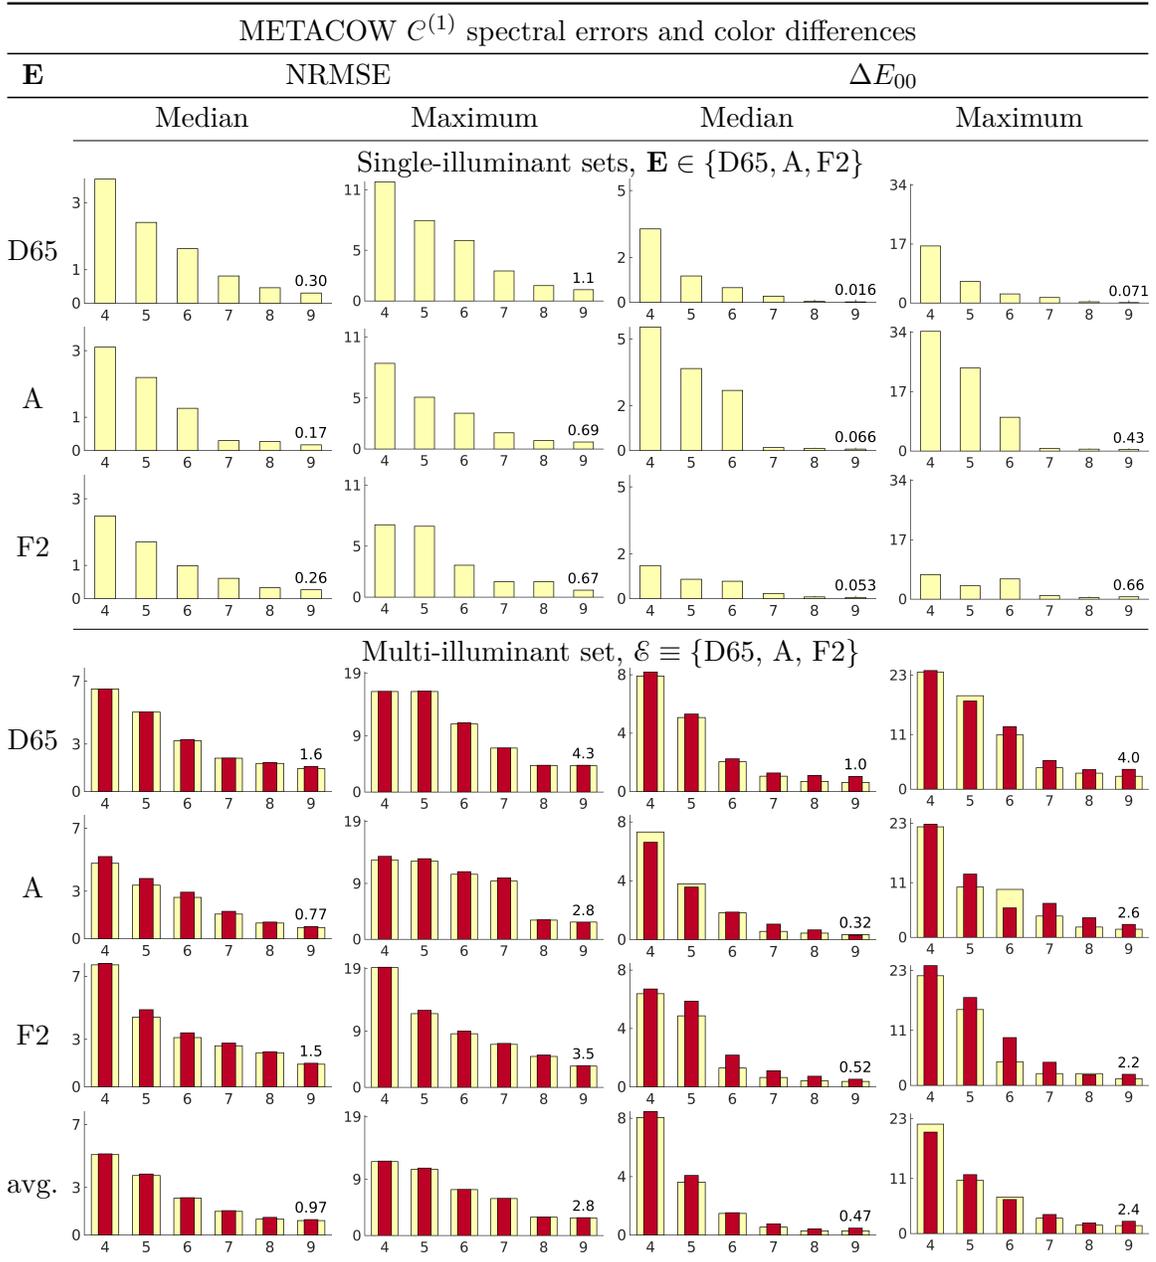


Figure 6.3: For dimensionalities $m \equiv 4, \dots, 9$, the spectral errors and color differences incurred when computing METACOW $\mathcal{C}^{(1)}$ reflections of D65, A, F2, or all three. In the latter case, errors incurred by reflecting the average all three illuminants are shown in the bottom row. For single-illuminant sets, reflections are computed by multiplying coefficients of sharp basis $\mathbf{Q}^{(\mathbf{V})} \equiv \mathbf{B}^{(1)}\mathbf{T}^{(\mathbf{V})}$, where change of basis matrices $\mathbf{T}^{(\mathbf{V})}$ were found by diagonalizing $\mathbf{B}^{(1)\text{T}} \text{diag}(\mathbf{E})^{-1}\mathbf{B}^{(1)}$, as described in Section 5.1. The inner (darker) bars of the multi-illuminant set represent errors incurred using sharp basis $\mathbf{Q}^{(1)} \equiv \mathbf{B}^{(1)}\mathbf{T}^{(1)}$, where $\mathbf{T}^{(1)}$ minimizes $f^{(1)}$, the total approximation error, as described in Section 5.5. Outer (lighter) bars represent errors incurred by projection onto $\mathbf{B}^{(1)}$.

6.1.1 Single light source, direct illumination only ($q \equiv k \equiv 1$)

Median and maximum spectral errors and color differences incurred for single- and multi-illuminant $\mathcal{C}^{(1)}$ SPDs are shown in Figs. 6.1-6.3. Errors incurred when computing reflections of a single light source SPD $\mathbf{E} \in \{\text{D65, A, F2}\}$ by multiplying coefficients of $\mathbf{Q}^{(\mathbf{V})}$, found by diagonalizing $\mathbf{B}^{(1)\text{T}} \text{diag}(\mathbf{E})^{-1} \mathbf{B}^{(1)}$, as described in Section 5.1, result in maximum color differences that would likely be imperceptible using, in some cases, as few as $m \equiv 5$ dimensions and $m \leq 9$ in all cases. These errors are equal to those attained by projection onto $\mathbf{B}^{(1)}$ and cannot be improved.

Sets of $\mathcal{C}^{(1)}$ SPDs generated by all three light sources require approximately twice the number of dimensions to attain the same accuracy when computed by multiplying coefficients of $\mathbf{Q}^{(1)} \equiv \mathbf{B}^{(1)} \mathbf{T}^{(1)}$ found by minimizing $f^{(1)}$, the total approximation error, as described in Section 5.5, but are still well approximated with nine or fewer, yielding median color differences of 1 or less. Some outlying maximum errors, causing color differences as high as 4.5, might be noticeable but tolerable, depending on the image. (Detailed distributions of spectral errors and color differences incurred for $\mathcal{C}^{(1)}$ SPDs, shown as histograms in Figs. A1-A12, are included in the appendix as supplementary material.)

All worst $\mathcal{C}^{(1)}$ approximations are depicted in Figs. 6.4-6.11, along with, for comparison, those incurred using $\mathbf{Q}^{(5)} \equiv \mathbf{B}^{(1)} \mathbf{T}^{(5)}$, found by maximizing sharpness by minimizing $f^{(5)}$, as described in Section 5.4. Figs. 6.1-6.3 and Fig. 6.10 show that almost all of the approximation error results from projection onto $\mathbf{B}^{(1)}$. That is, very little additional error is incurred by computing reflections as the product of coefficients of sharp basis $\mathbf{Q}^{(1)}$.

6.1.2 Direct and indirect illumination, $k > 1$

Fig. 6.12 shows results obtained for single-illuminant Munsell $\mathcal{C}^{(3)}$ SPDs using bases $\mathbf{Q}^{(i)} \equiv \mathbf{B}^{(2)} \mathbf{T}^{(i)}$ found by sharpening alone, as described in Section 5.4, where $\mathbf{T}^{(i)}$ optimizes $f^{(i)} \in \{f^{(4)}, f^{(5)}, f^{(6)}\}$, i.e. the method of Drew and Finlayson, quartimin or quartimax. (Multi-illuminant sets, as well as Natural Colors and METACOW sets, yield similar results.) Also included, for comparison, are results obtained using $\mathbf{Q}^{(1)} \equiv \mathbf{B}^{(2)} \mathbf{T}^{(1)}$, found by explicit minimization of $f^{(1)}$, the total approximation error, as well as the lower bound errors incurred by projection onto $\mathbf{B}^{(2)}$. The graphs in the last row of this figure, in which the sharpness of each transformed basis is plotted (in negative log scale for clarity), shows that loosening the quartimax orthogonality constraint on

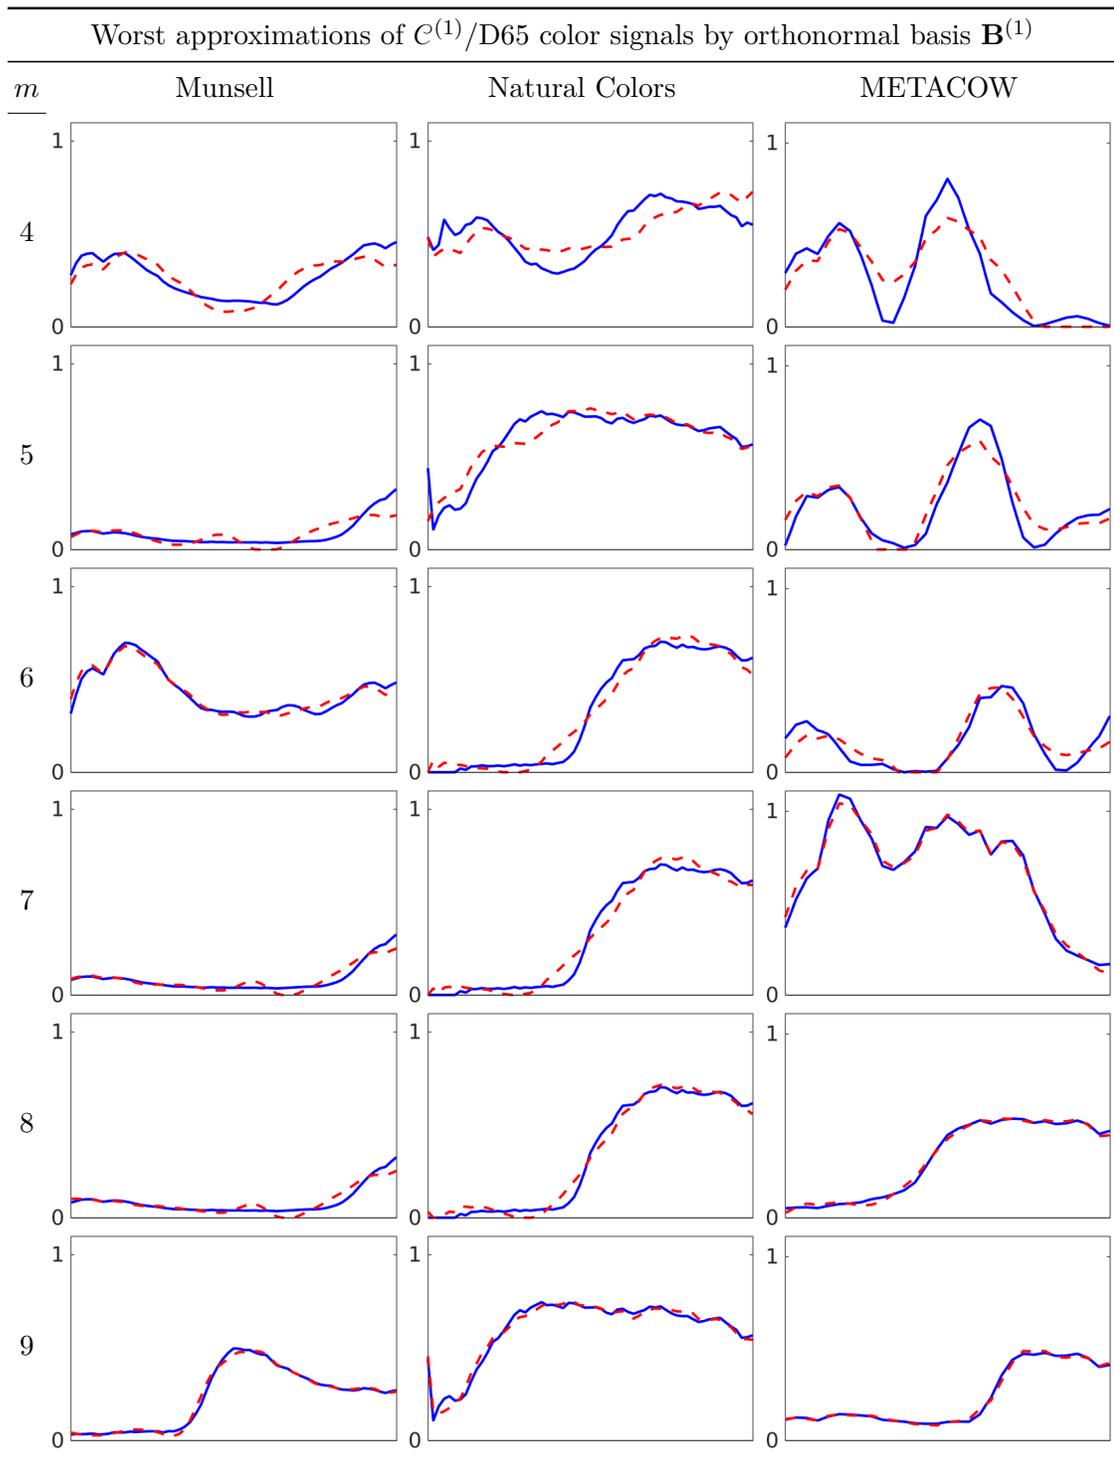


Figure 6.4: From $\mathcal{C}^{(1)}$ sets constructed with D65, exact color signals are plotted with blue (solid) lines. For dimensionalities $m \equiv 4, \dots, 9$, projections onto $\mathbf{B}^{(1)}$ are plotted with red (dashed) lines.

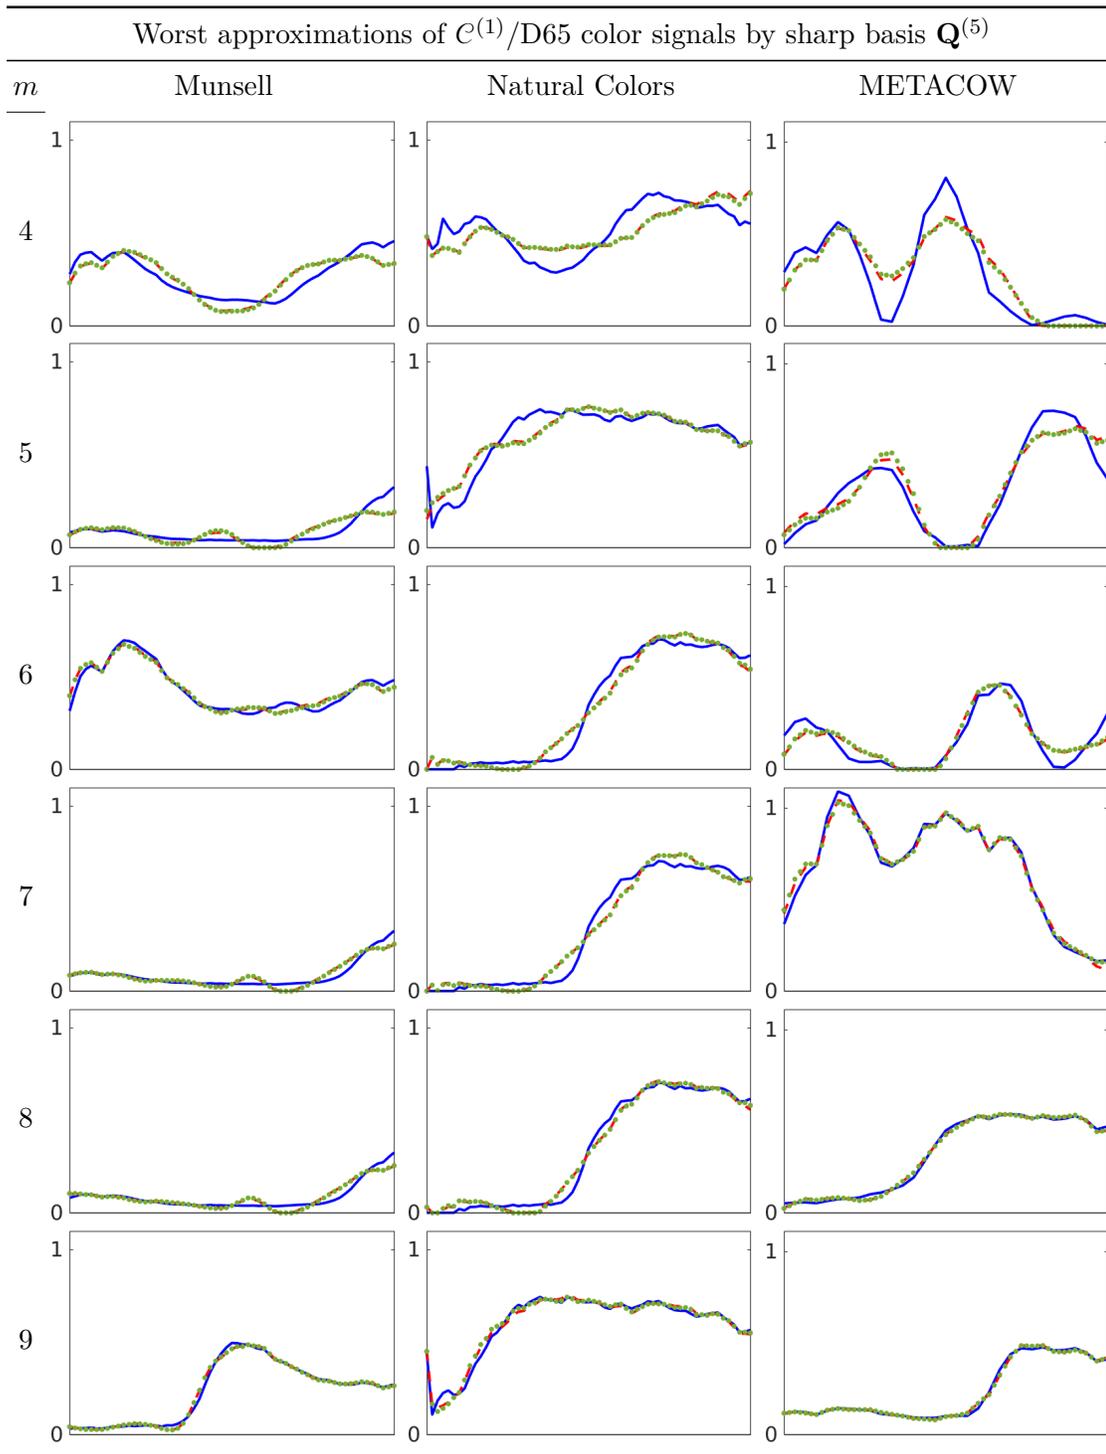


Figure 6.5: From $\mathcal{C}^{(1)}$ sets constructed with D65, exact color signals are plotted with blue (solid) lines. For dimensionalities $m \equiv 4, \dots, 9$, approximations computed by multiplying coefficients of sharp bases $\mathbf{Q}^{(5)} \equiv \mathbf{B}^{(1)}\mathbf{T}^{(5)}$ are plotted with green (dotted) lines. For comparison, projections onto $\mathbf{B}^{(1)}$ are plotted with red (dashed) lines.

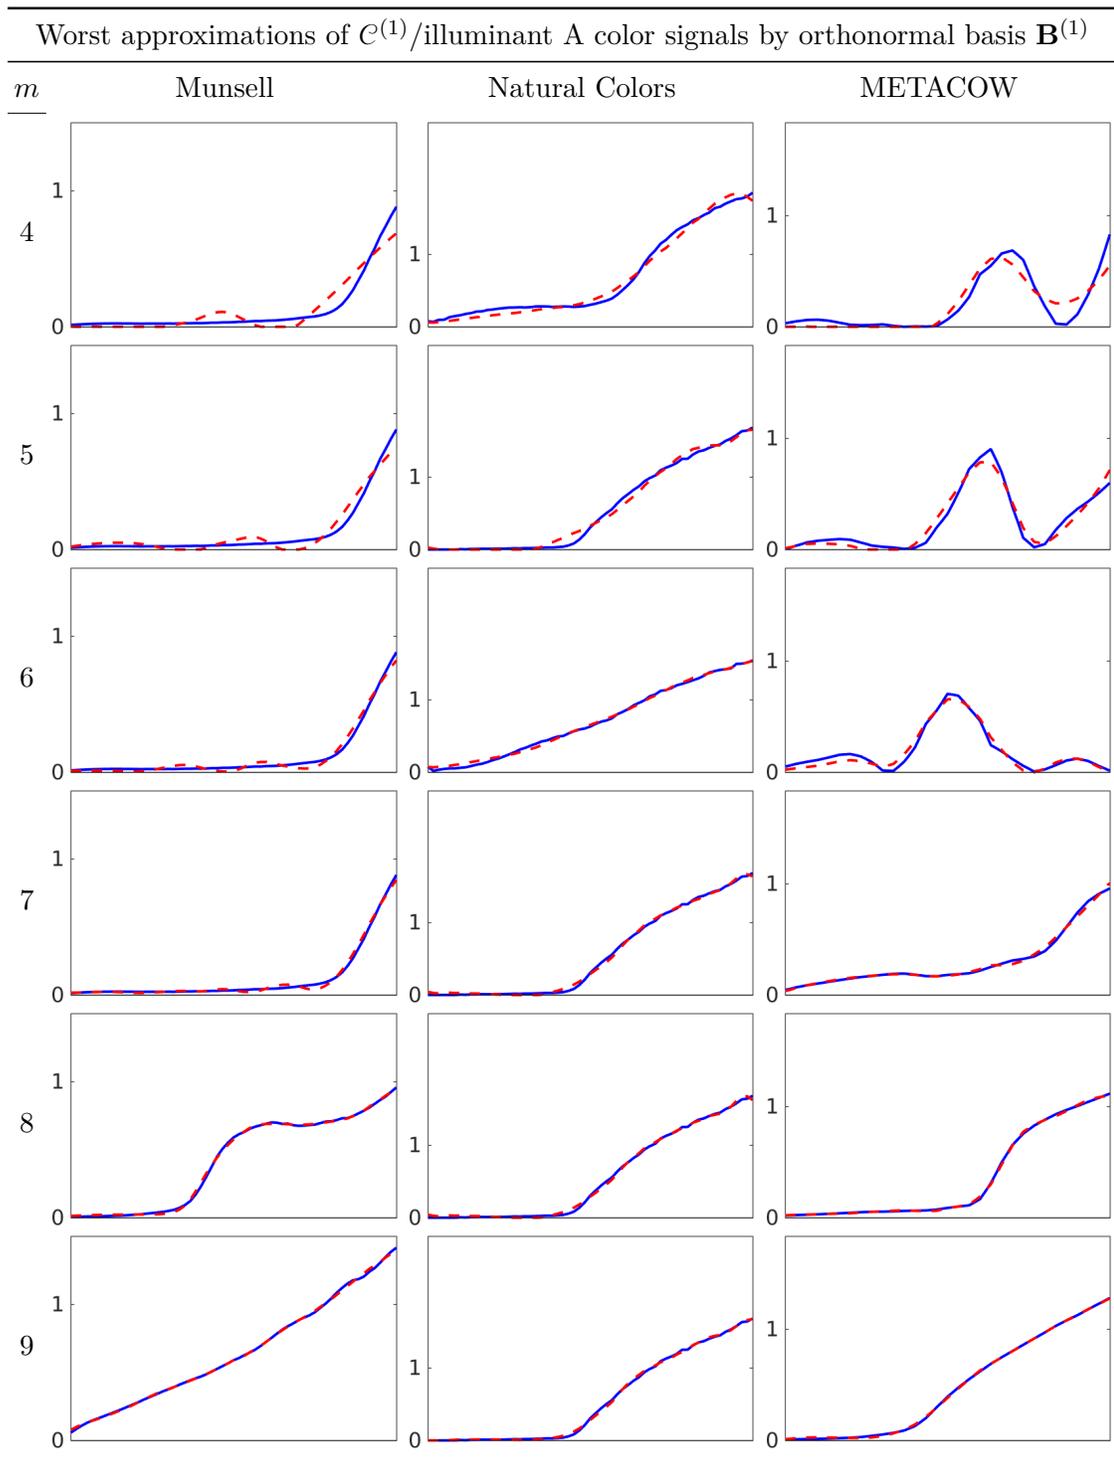


Figure 6.6: From $\mathcal{C}^{(1)}$ sets constructed with illuminant A, exact color signals are plotted with blue (solid) lines. For dimensionalities $m \equiv 4, \dots, 9$, projections onto $\mathbf{B}^{(1)}$ are plotted with red (dashed) lines.

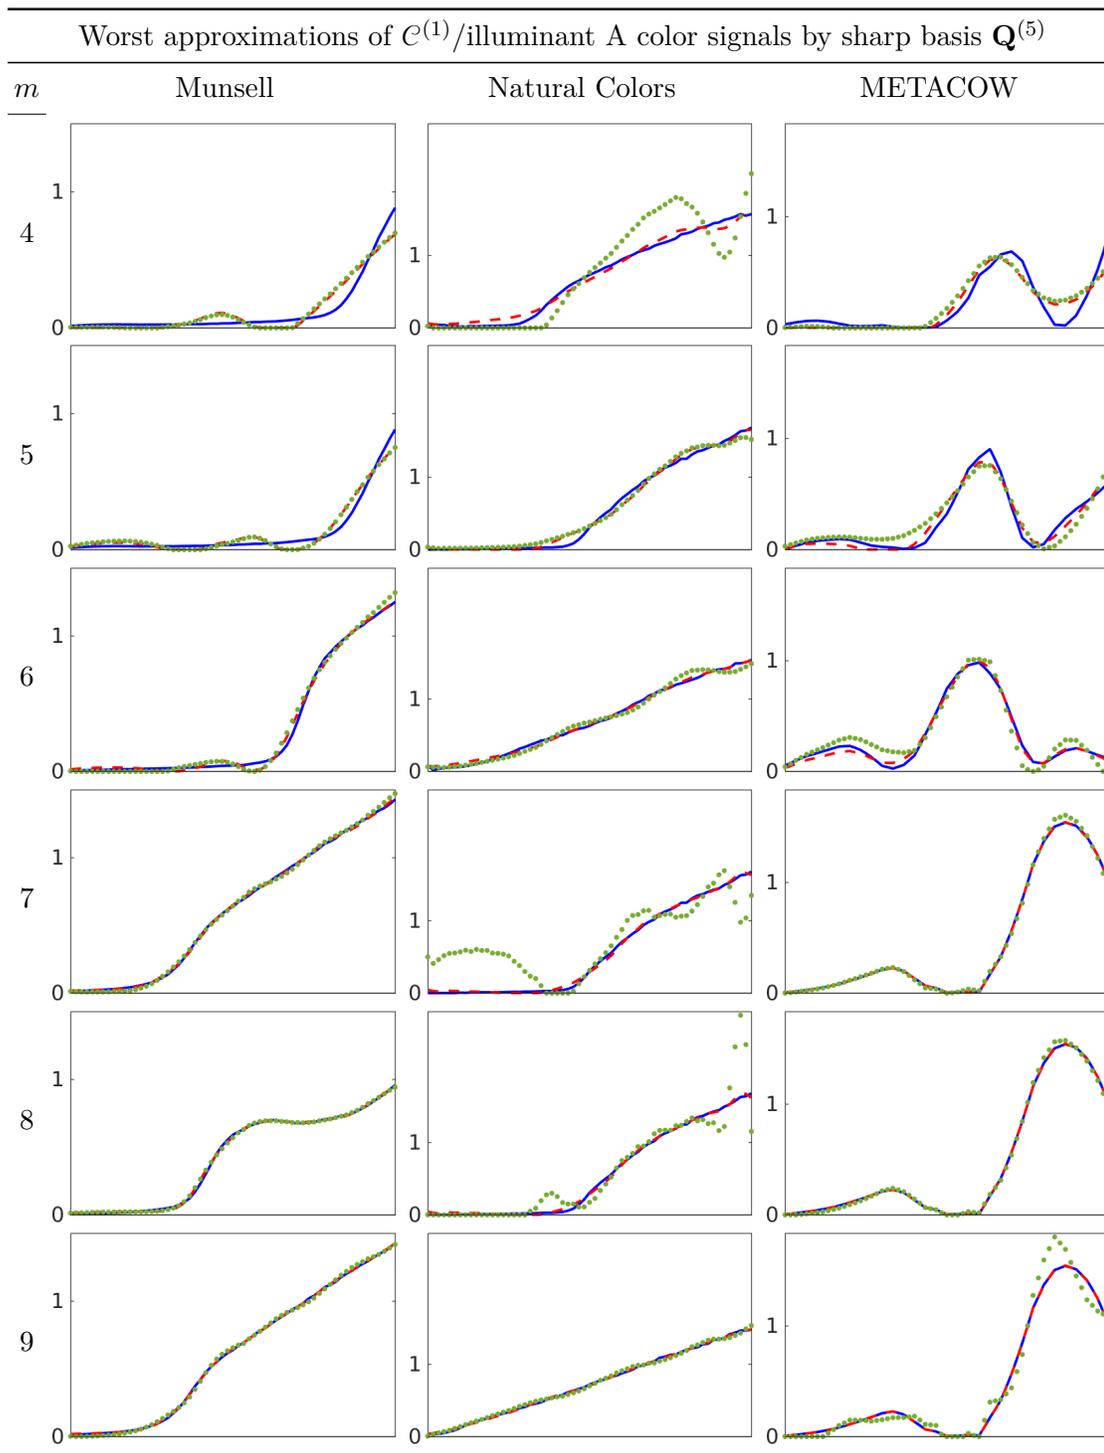


Figure 6.7: From $\mathcal{C}^{(1)}$ sets constructed with illuminant A, exact color signals are plotted with blue (solid) lines. For dimensionalities $m \equiv 4, \dots, 9$, approximations computed by multiplying coefficients of sharp bases $\mathbf{Q}^{(5)} \equiv \mathbf{B}^{(1)}\mathbf{T}^{(5)}$ are plotted with green (dotted) lines. For comparison, projections onto $\mathbf{B}^{(1)}$ are plotted with red (dashed) lines.

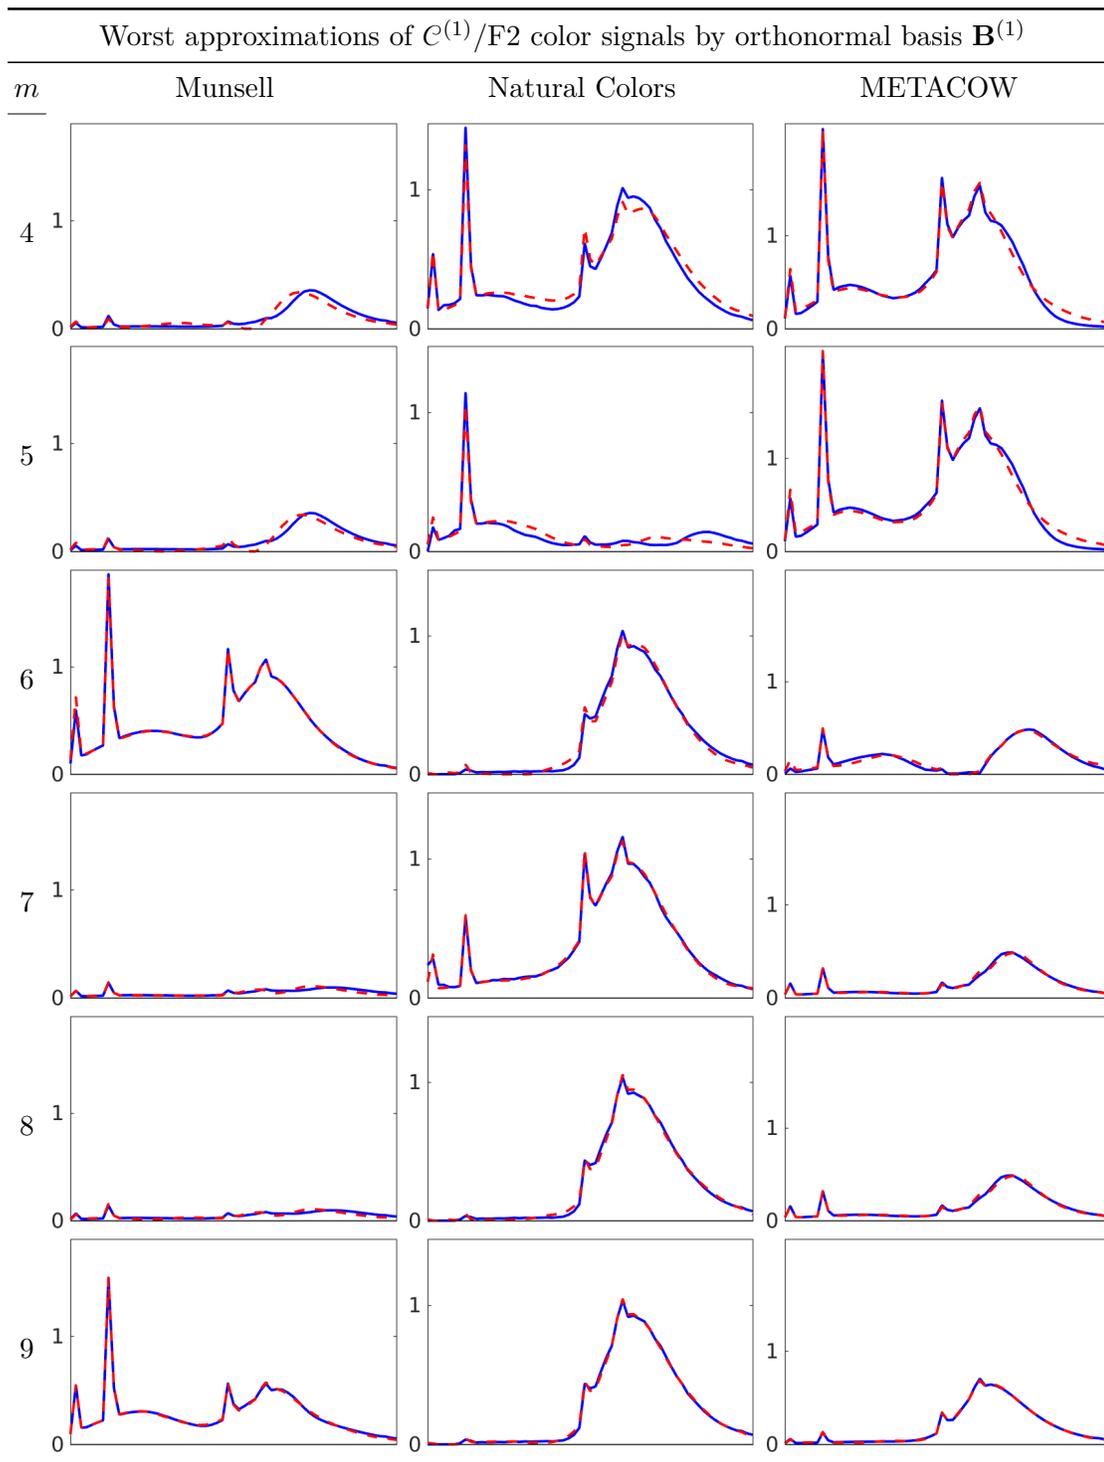


Figure 6.8: From $\mathcal{C}^{(1)}$ sets constructed with illuminant F2, exact color signals are plotted with blue (solid) lines. For dimensionalities $m \equiv 4, \dots, 9$, projections onto $\mathbf{B}^{(1)}$ are plotted with red (dashed) lines.

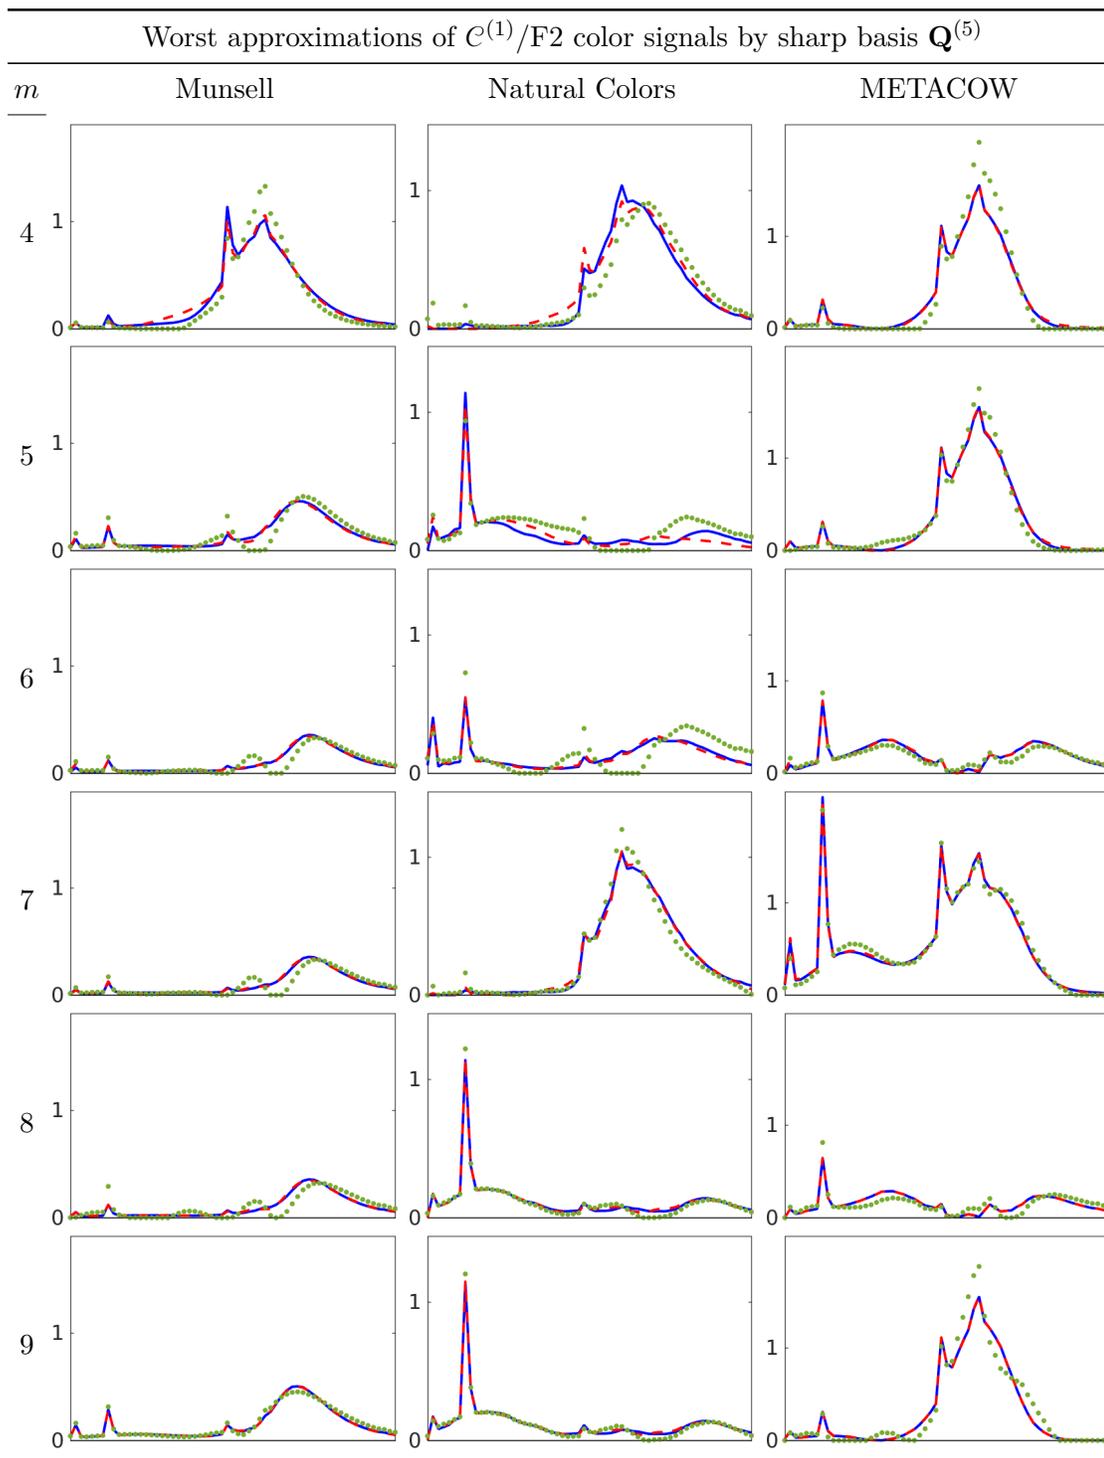


Figure 6.9: From $\mathcal{C}^{(1)}$ sets constructed with illuminant F2, exact color signals are plotted with blue (solid) lines. For dimensionalities $m \equiv 4, \dots, 9$, approximations computed by multiplying coefficients of sharp bases $\mathbf{Q}^{(5)} \equiv \mathbf{B}^{(1)}\mathbf{T}^{(5)}$ are plotted with green (dotted) lines. For comparison, projections onto $\mathbf{B}^{(1)}$ are plotted with red (dashed) lines.

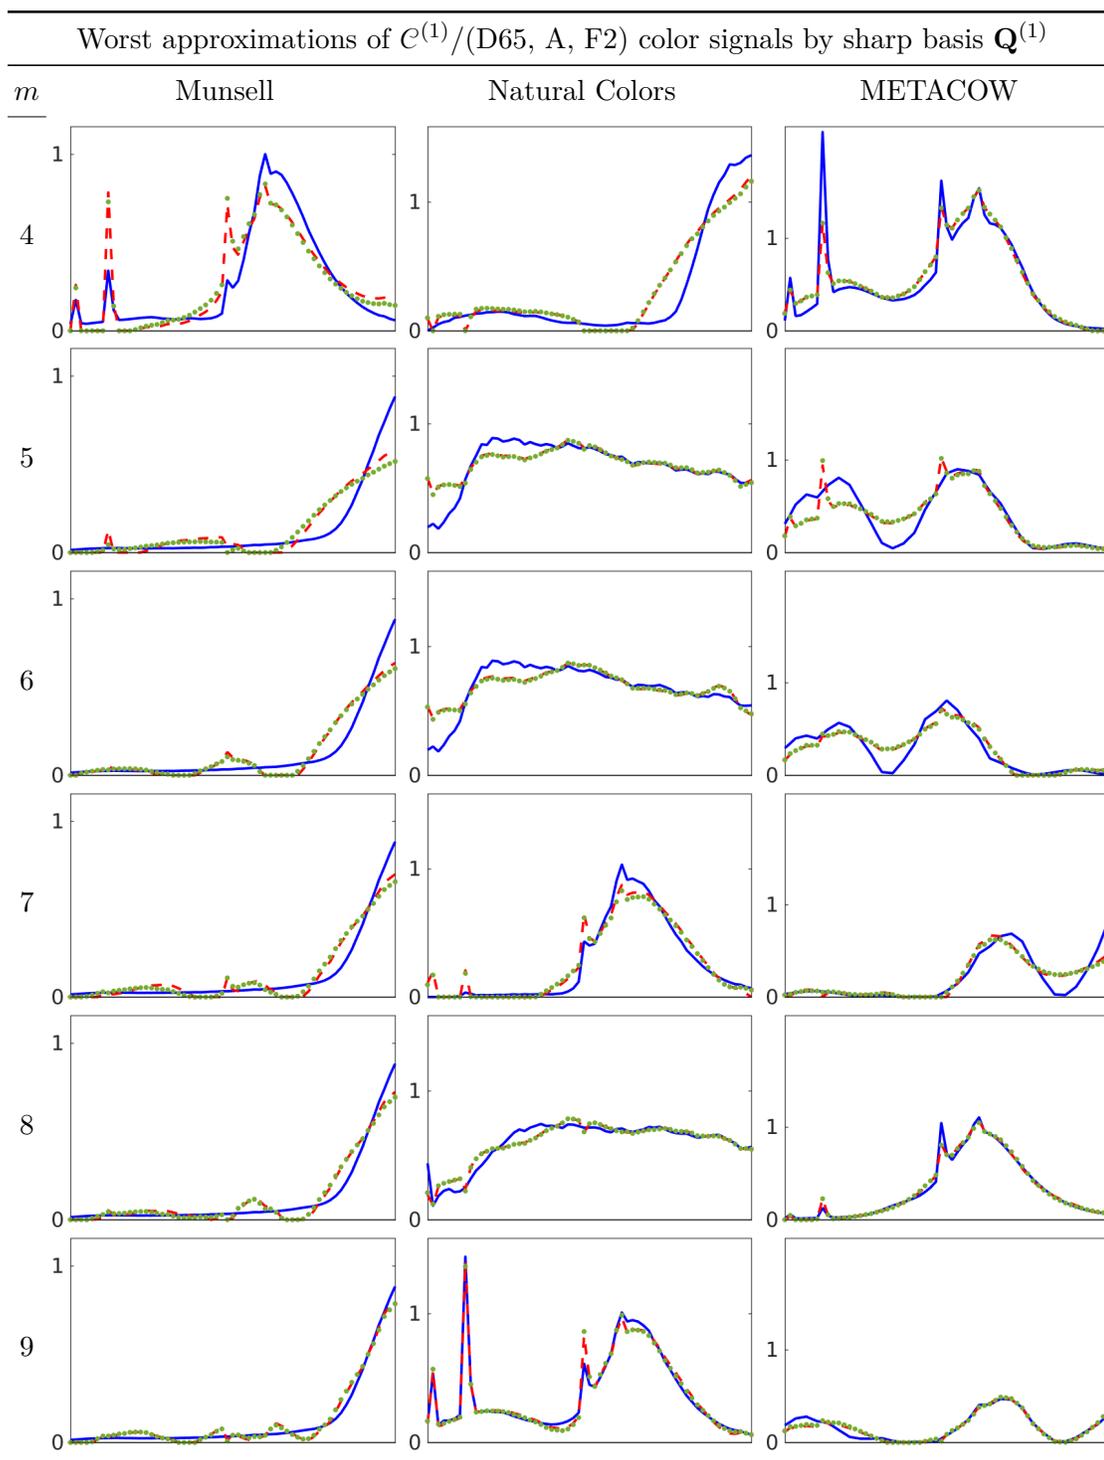


Figure 6.10: From $\mathcal{C}^{(1)}$ sets constructed with illuminants D65, A and F2, exact color signals are plotted with blue (solid) lines. For dimensionalities $m \equiv 4, \dots, 9$, approximations computed by multiplying coefficients of sharp bases $\mathbf{Q}^{(1)} \equiv \mathbf{B}^{(1)}\mathbf{T}^{(1)}$ are plotted with green (dotted) lines. For comparison, projections onto $\mathbf{B}^{(1)}$ are plotted with red (dashed) lines.

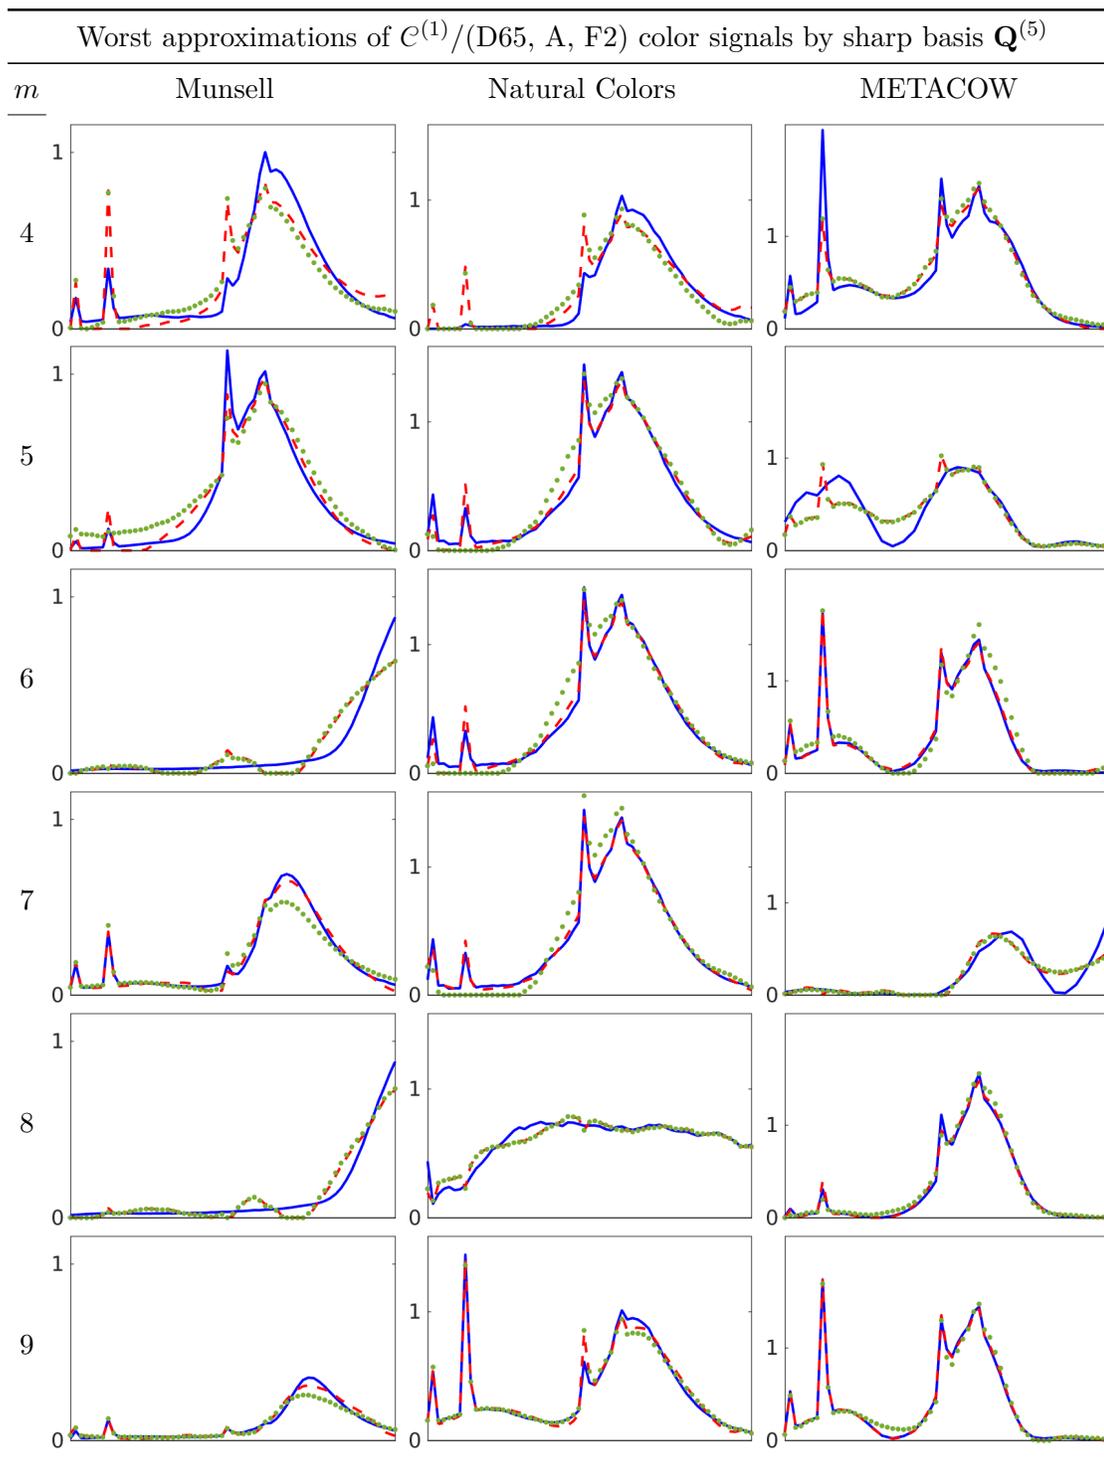


Figure 6.11: From $\mathcal{C}^{(1)}$ sets constructed with illuminants D65, A and F2, exact color signals are plotted with blue (solid) lines. For dimensionalities $m \equiv 4, \dots, 9$, approximations computed by multiplying coefficients of sharp bases $\mathbf{Q}^{(5)} \equiv \mathbf{B}^{(1)}\mathbf{T}^{(5)}$ are plotted with green (dotted) lines. For comparison, projections onto $\mathbf{B}^{(1)}$ are plotted with red (dashed) lines.

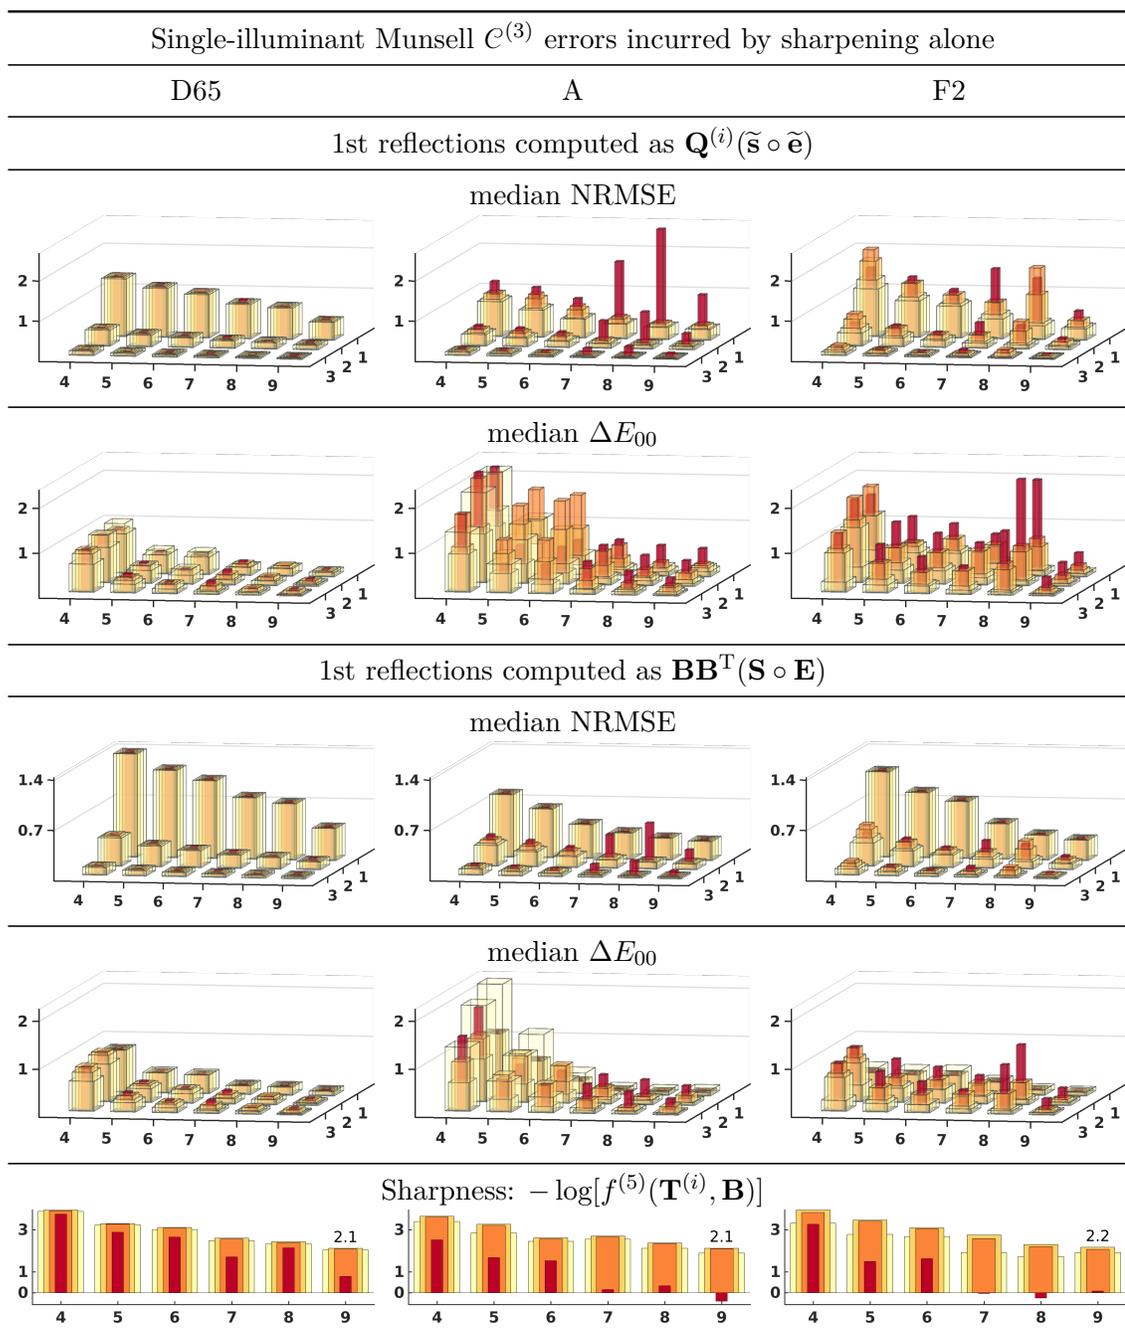


Figure 6.12: For dimensionalities $m \equiv 4, \dots, 9$, and 1st, 2nd and 3rd reflections, the spectral errors and color differences incurred when computing Munsell $\mathcal{C}^{(3)}$ reflections of D65, A, or F2 by multiplying coefficients of $\mathbf{Q}^{(i)} \equiv \mathbf{B}^{(2)}\mathbf{T}^{(i)}$. Inner (darker) to outer (lighter) nested 3D bars represent errors incurred using, respectively, $\mathbf{Q}^{(4)}$, $\mathbf{Q}^{(6)}$, $\mathbf{Q}^{(5)}$ or $\mathbf{Q}^{(1)}$ and by projection onto $\mathbf{B}^{(2)}$. Bars in bottom row represent the sharpness of $\mathbf{Q}^{(i)}$ in the same order, plotted in negative log scale for clarity (higher bars mean sharper bases).

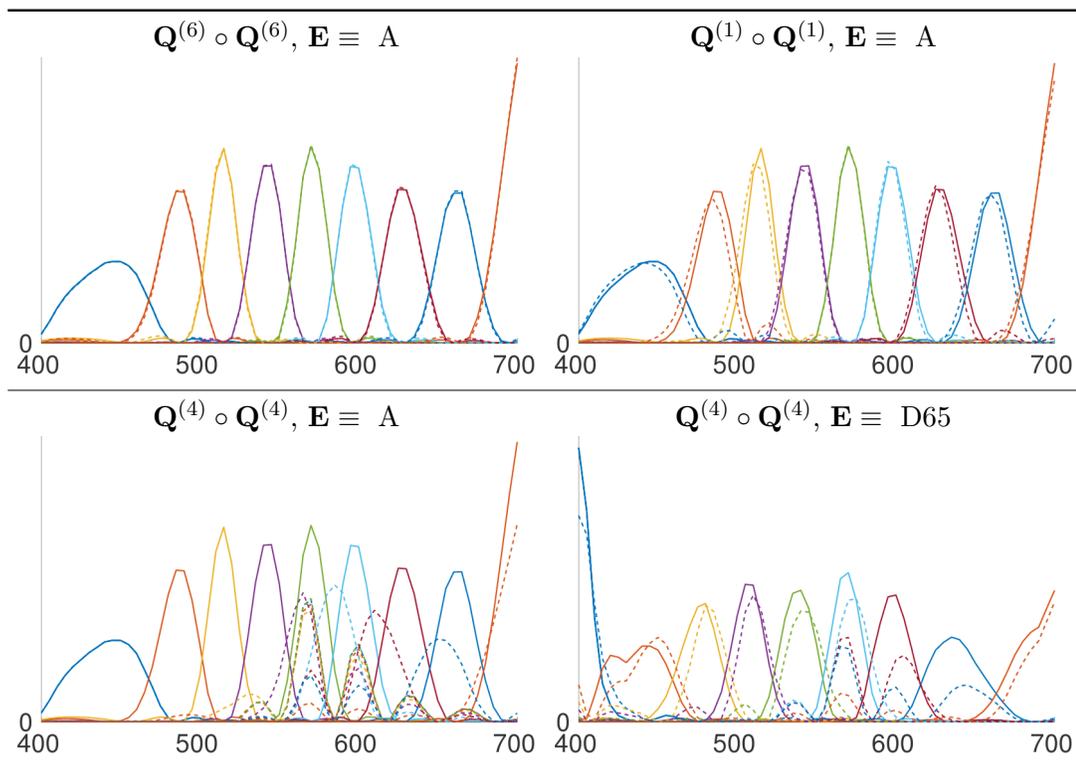


Figure 6.13: Squared nine-dimensional bases $\mathbf{Q}^{(i)} \equiv \mathbf{B}^{(2)}\mathbf{T}^{(i)}$, for Munsell SPDs formed with illuminant A or D65, where $\mathbf{T}^{(i)}$ minimizes $f^{(i)}$, $i \in \{6, 1, 4\}$. To compare sharpness, maximally sharp $\mathbf{Q}^{(5)} \circ \mathbf{Q}^{(5)}$ is plotted in each graph (with solid lines) along with $\mathbf{Q}^{(i)} \circ \mathbf{Q}^{(i)}$ (with dashed lines). The sharpness of these bases, and the approximation errors incurred when reflections are computed by multiplying their coefficients, are all plotted in Fig. 6.12.

\mathbf{T} to quartimin's unit length constraint on its columns, increases the sharpness of $\mathbf{Q}^{(5)}$ by only a small amount. In general, $\mathbf{Q}^{(5)} \approx \mathbf{Q}^{(6)}$, as shown by a typical example in Fig. 6.13. The sharpness and error graphs in Fig. 6.12 also show, for these examples, that while the various $\mathbf{Q}^{(1)}$ yield, as expected, more accurate approximations than the corresponding $\mathbf{Q}^{(5)}$ and $\mathbf{Q}^{(6)}$, they are consistently less sharp, which is also expected. However, these results also show that accuracy is not unrelated to sharpness, as all $\mathbf{Q}^{(1)}$ are almost as sharp. Conversely, the $\mathbf{Q}^{(4)}$ bases obtained by the method of Drew and Finlayson, using approximately uniform, evenly distributed and connected sharpening intervals $\{\mathcal{L}^{(i)}\}_{i=1}^m$, are not only significantly and consistently less sharp, but they also yield in most cases the least accurate approximations.

Approximations of $\mathcal{C}^{(1)}$ SPDs by their projections onto the basis $\mathbf{B}^{(2)}$ for the corresponding $\mathcal{C}^{(2)}$ SPDs are necessarily less accurate than their projections onto the $\mathbf{B}^{(1)}$ basis of the same dimensionality. In general, because the leading vectors

of $\mathbf{B}^{(1)}$ and $\mathbf{B}^{(2)}$ tend to be similar, the difference is not large and likely tolerable, but, if necessary, equivalent accuracy can be attained with a $\mathbf{B}^{(2)}$ basis at the cost of increasing its dimensionality by one, or possibly two in some cases. More concerning is that the accuracy attained with a basis found by sharpening alone cannot always be improved by adding a dimension. In some cases, as shown in Fig. 6.12, the error incurred by reflections of illuminant A or F2 increases with higher dimensionalities. That this does not occur with $\mathbf{Q}^{(1)}$ bases, as is also shown by the examples in Fig. 6.12, is presumably a result of the rapidly decreasing singular values of $\mathbf{c} \equiv \mathbf{B}^T \mathcal{C}^{k-1}$, which, as their number increases with m , provide more degrees of freedom to orient $\mathbf{V}^{(\Delta)}$ with respect to $\mathbf{U}^{(\mathcal{C})}$ while increasing $\|\Delta\|$ without increasing $\|\Delta\mathbf{c}\|$.

Small differences between $\mathbf{Q}^{(1)}$ and $\mathbf{Q}^{(5)}$ or $\mathbf{Q}^{(6)}$ can have a large effect on the approximation error incurred by reflections of illuminant A or F2. For example, Figs. 6.12 and 6.13 show that although, for Munsell/illuminant A SPDs, $\mathbf{Q}^{(1)} \approx \mathbf{Q}^{(5)}$, approximately 32%, 46% and 49% greater total spectral error is incurred by $\mathbf{Q}^{(5)}$ for first, second and third reflections, respectively. For D65 reflections, on the other hand, $\mathbf{Q}^{(1)}$, $\mathbf{Q}^{(4)}$, $\mathbf{Q}^{(5)}$ and $\mathbf{Q}^{(6)}$ bases all attain approximately the highest possible accuracy, although they can be quite different, as shown, for example, by the nine-dimensional $\mathbf{Q}^{(4)}$ and $\mathbf{Q}^{(5)}$ in Fig. 6.13. This equivalence results from the normalization $\mathbf{Q}^+ \mathbf{1} = \mathbf{1}$, and the small variation in D65's SPD, which is sufficiently close to a constant function. Observe that because

$$\|\Delta\mathbf{c}\| = \|\Sigma^{(\Delta)} \mathbf{V}^{(\Delta)T} \mathbf{U}^{(\mathcal{C})} \Sigma^{(\mathcal{C})}\| = \|\Sigma^{(\Delta)} \mathbf{V}^{(\Delta)T} \mathbf{B}^T \mathbf{U}^{(\mathcal{C})} \Sigma^{(\mathcal{C})}\|, \quad (6.1)$$

where $\mathcal{C} \equiv \mathcal{C}^{(1)} = \mathbf{U}^{(\mathcal{C})} \Sigma^{(\mathcal{C})} \mathbf{V}^{(\mathcal{C})T}$, when an upper bound on $\|\Delta\|$ is implicitly reduced by spectral sharpening, $\|\Delta\mathbf{c}\|$ will be small when three conditions are met: 1) most of $\text{tr}(\Sigma^{(\mathcal{C})})$ is in a few leading dimensions, concentrated mostly in the leading singular value $\sigma_1^{(\mathcal{C})}$; 2) the trailing singular value $\sigma_m^{(\Delta)} \approx 0$; and 3) the rightmost columns of $\mathbf{B}\mathbf{V}^{(\Delta)}$ are approximately parallel to the leftmost columns of $\mathbf{U}^{(\mathcal{C})}$, in particular $\langle \mathbf{B}\mathbf{V}_m^{(\Delta)}, \mathbf{U}_1^{(\mathcal{C})} \rangle \approx 1$.

The first of these conditions holds in general as a consequence of the high correlation of $\mathcal{C}^{(1)}$ SPDs. To understand why the second and third conditions hold when $\mathbf{Q}^+ \mathbf{1} = \mathbf{1}$ and $\mathbf{B}\mathbf{B}^T \mathbf{1} \approx \mathbf{1}$, observe, for an SPD $\mathbf{C} \equiv \alpha \mathbf{1} \approx \alpha \mathbf{B}\mathbf{B}^T \mathbf{1}$, for some constant

α , that

$$\begin{aligned}\Delta^{(i)}\mathbf{B}^T\mathbf{C} &= \alpha[\mathbf{B}^T \text{diag}(\widehat{\mathbf{U}}_i^{(s)})\mathbf{B} - \mathbf{T} \text{diag}(\widetilde{\mathbf{u}}_i^{(s)})\mathbf{T}^{-1}]\mathbf{B}^T\mathbf{1} \\ &\approx \alpha[\mathbf{B}^T \text{diag}(\widehat{\mathbf{U}}_i^{(s)})\mathbf{1} - \mathbf{T} \text{diag}(\widetilde{\mathbf{u}}_i^{(s)})\mathbf{1}] = \alpha(\mathbf{B}^T\widehat{\mathbf{U}}_i^{(s)} - \mathbf{T}\widetilde{\mathbf{u}}_i^{(s)}) = \mathbf{0}.\end{aligned}$$

Thus, because $\|\Delta\mathbf{c} - \mathbf{c}\|$ is smallest in directions \mathbf{c} near a constant SPD's coordinates, $\mathbf{V}_m^{(\Delta)} \approx (1, 0, \dots, 0)^T$, if $\langle \mathbf{B}_1, \mathbf{1} \rangle \approx \|\mathbf{1}\|$. This condition holds sufficiently for D65 $\mathcal{C}^{(2)}$ SPDs, as \mathbf{B}_1 , the leading characteristic vector of $\mathcal{C}^{(2)}$, is approximately the normalized mean $\mathcal{C}^{(2)}$ SPD. For well distributed reflectance spectra (i.e., a *gray world*), the mean approximates D65, which is sufficiently close to a constant function. As well, because the leading dimensions of $\mathbf{B}^{(1)}$ and $\mathbf{B}^{(2)}$ are approximately equal, $\mathbf{B}_{1:i}^{(2)T}\mathbf{U}_{1:i}^{(e)} \approx \mathbf{I}$, for $i \approx 2, 3$ or higher. In particular, $\mathbf{B}_1^{(2)} \approx \mathbf{B}_1^{(1)} = \mathbf{U}_1^{(e)}$. Furthermore, $\mathbf{V}^{(\Delta)}$ will tend to approximate, at least in the trailing dimensions, an anti-diagonal matrix with ones along the anti-diagonal, since it will usually hold that $\|\Delta\mathbf{B}^T\mathbf{U}_i^{(e)}\| < \|\Delta\mathbf{B}^T\mathbf{U}_{i+1}^{(e)}\|$ for at least the first few columns of $\mathbf{U}^{(e)}$, which approximate sinusoids of increasing frequency.

An heuristic explanation of this observation can be made by noting that

$$\begin{aligned}\|\Delta^{(i)}\mathbf{B}^T\mathbf{U}_j^{(e)}\| &= \|\mathbf{B}\Delta^{(i)}\mathbf{B}^T\mathbf{U}_j^{(e)}\| \\ &= \|\mathbf{B}\mathbf{B}^T \text{diag}(\widehat{\mathbf{U}}_i^{(s)})\mathbf{B}\mathbf{B}^T\mathbf{U}_j^{(e)} - \mathbf{Q} \text{diag}(\mathbf{Q}^+\widehat{\mathbf{U}}_i^{(s)})\mathbf{Q}^+\mathbf{U}_j^{(e)}\| \\ &\approx \|\widehat{\mathbf{U}}_i^{(s)} \circ \mathbf{U}_j^{(e)} - \mathbf{Q} \text{diag}(\mathbf{Q}^+\widehat{\mathbf{U}}_i^{(s)})\mathbf{Q}^+\mathbf{U}_j^{(e)}\|,\end{aligned}\tag{6.2}$$

where the approximation holds for low values i and j , but not for $i, j > m$. The difference in (6.2) is a measure of the accuracy of approximating the product of two quasi-sinusoids by the product of their approximate average values in the ranges spanned by each roughly disjoint column of \mathbf{Q}^+ . Since the accuracy of this approximation, which is similar to a box filter reconstruction, decreases with increasing frequency of $\mathbf{U}_j^{(e)}$, the right singular vectors of Δ will tend to approximate, at least in the trailing dimensions, an anti-diagonal matrix with ones along the anti-diagonal. This favorable orientation of $\mathbf{V}^{(\Delta)}$ with respect to $\mathbf{B}^T\mathbf{U}^{(e)}$, obtained explicitly by minimization of $f^{(1)}$ and, for D65 $\mathcal{C}^{(2)}$ SPDs, implicitly by simple spectral sharpening alone, reduces the weighting of Δ by the leading singular values of \mathcal{C} , where it matters most. As a concrete example, for Munsell D65 reflections, $\mathbf{T}^{(5)}$ and $m \equiv 4$

dimensions, displayed rounded to the first decimal place,

$$\begin{aligned}
& \Sigma^{(\Delta)} \mathbf{V}^{(\Delta)\text{T}} [\mathbf{B}^{\text{T}} \mathbf{U}_{1:m}^{(c)}] \Sigma_{1:m}^{(c)} \\
&= \text{diag} \begin{bmatrix} 1.2 \\ 1.1 \\ 0.4 \\ 0.1 \end{bmatrix} \begin{bmatrix} 0.0 & -0.3 & -0.6 & 0.7 \\ 0.0 & -0.3 & 0.8 & 0.5 \\ 0.0 & 0.9 & 0.1 & 0.4 \\ 1.0 & 0.0 & 0.0 & 0.0 \end{bmatrix} \begin{bmatrix} 1.0 & -0.1 & 0.0 & 0.0 \\ 0.1 & 1.0 & -0.1 & 0.0 \\ 0.0 & 0.1 & 1.0 & -0.1 \\ 0.0 & 0.0 & 0.1 & 1.0 \end{bmatrix} \text{diag} \begin{bmatrix} 81.8 \\ 19.6 \\ 12.5 \\ 4.4 \end{bmatrix} \\
&= \text{diag} \begin{bmatrix} 1.2 \\ 1.1 \\ 0.4 \\ 0.1 \end{bmatrix} \begin{bmatrix} 0.0 & -0.4 & -0.5 & 0.8 \\ 0.0 & -0.2 & 0.9 & 0.4 \\ 0.1 & 0.9 & 0.0 & 0.4 \\ 1.0 & -0.1 & 0.0 & 0.0 \end{bmatrix} \text{diag} \begin{bmatrix} 81.8 \\ 19.6 \\ 12.5 \\ 4.4 \end{bmatrix}.
\end{aligned}$$

Error incurred when computing a reflection with sharp basis coefficients is propagated and increased on subsequent reflections. While this is not a significant problem for D65 illumination, large errors can be incurred by first reflections of illuminant A or F2. As examples of the contribution of first reflection errors to subsequent reflections, the right two columns of Fig. 6.12 show the median errors incurred when first reflections are computed as projections onto \mathbf{B} and propagated by multiplying coefficients of \mathbf{Q} . For illuminant A and F2 Munsell SPDs and all dimensionalities from $m \equiv 4$ through 9, the first reflection errors incurred by $\mathbf{Q}^{(5)}$ and $\mathbf{Q}^{(1)}$ bases respectively contribute from 12% to 49% and from 1% to 10% to second reflection errors. The same errors respectively account for between 8% and 25% and between 1% and 6% of third reflection errors.

Figs. 6.14-6.16 shows results obtained for $\mathcal{C}^{(3)}$ SPDs by the four principal methods used to find \mathbf{T} described in the previous chapter: 1) spectral sharpening, by minimizing $f^{(5)}$ (Section 5.4); 2) AJD, by minimizing $f^{(3)}$ with the UWAJD algorithm (Section 5.3); 3) minimization of $f^{(2)} \equiv \|\Delta\|_F^2$ (Section 5.2); and 4) minimization of $f^{(1)} \equiv \|\Delta\mathbf{c}\|_F^2$, the total approximation error (Section 5.5). (Detailed distributions of $\mathcal{C}^{(3)}$ spectral errors and color differences, shown as histograms in Figs. A13-A36, are included in the appendix as supplementary material.) Also shown, for comparison, are the lower bound errors incurred by projection onto $\mathbf{B}^{(2)}$. Median errors incurred by all methods other than spectral sharpening are likely imperceptible when approximated with a basis of nine dimensions, or, in some cases, as few as four or five. Similar to the results shown in Fig. 6.12, we see that errors for D65 SPDs are relatively insensitive to the choice of \mathbf{T} , as all methods attain approximately the best lower bound. Similarly, as well, the error incurred by all methods other than

minimization of $f^{(1)}$ does not always decrease with increased dimensionality. The worst $\mathcal{C}^{(3)}$ approximations using $\mathbf{Q}^{(1)}$ are depicted in Figs. 6.17-6.24, along with, for comparison, those incurred using maximally sharp $\mathbf{Q}^{(5)}$. As is the case for the worst $\mathcal{C}^{(1)}$ approximations shown in Figs. 6.4-6.11, very little additional error is incurred by computing reflections as the product of coefficients of sharp basis $\mathbf{Q}^{(1)}$. That is, almost all of the approximation is due to the lower bound resulting from projection onto $\mathbf{B}^{(1)}$.

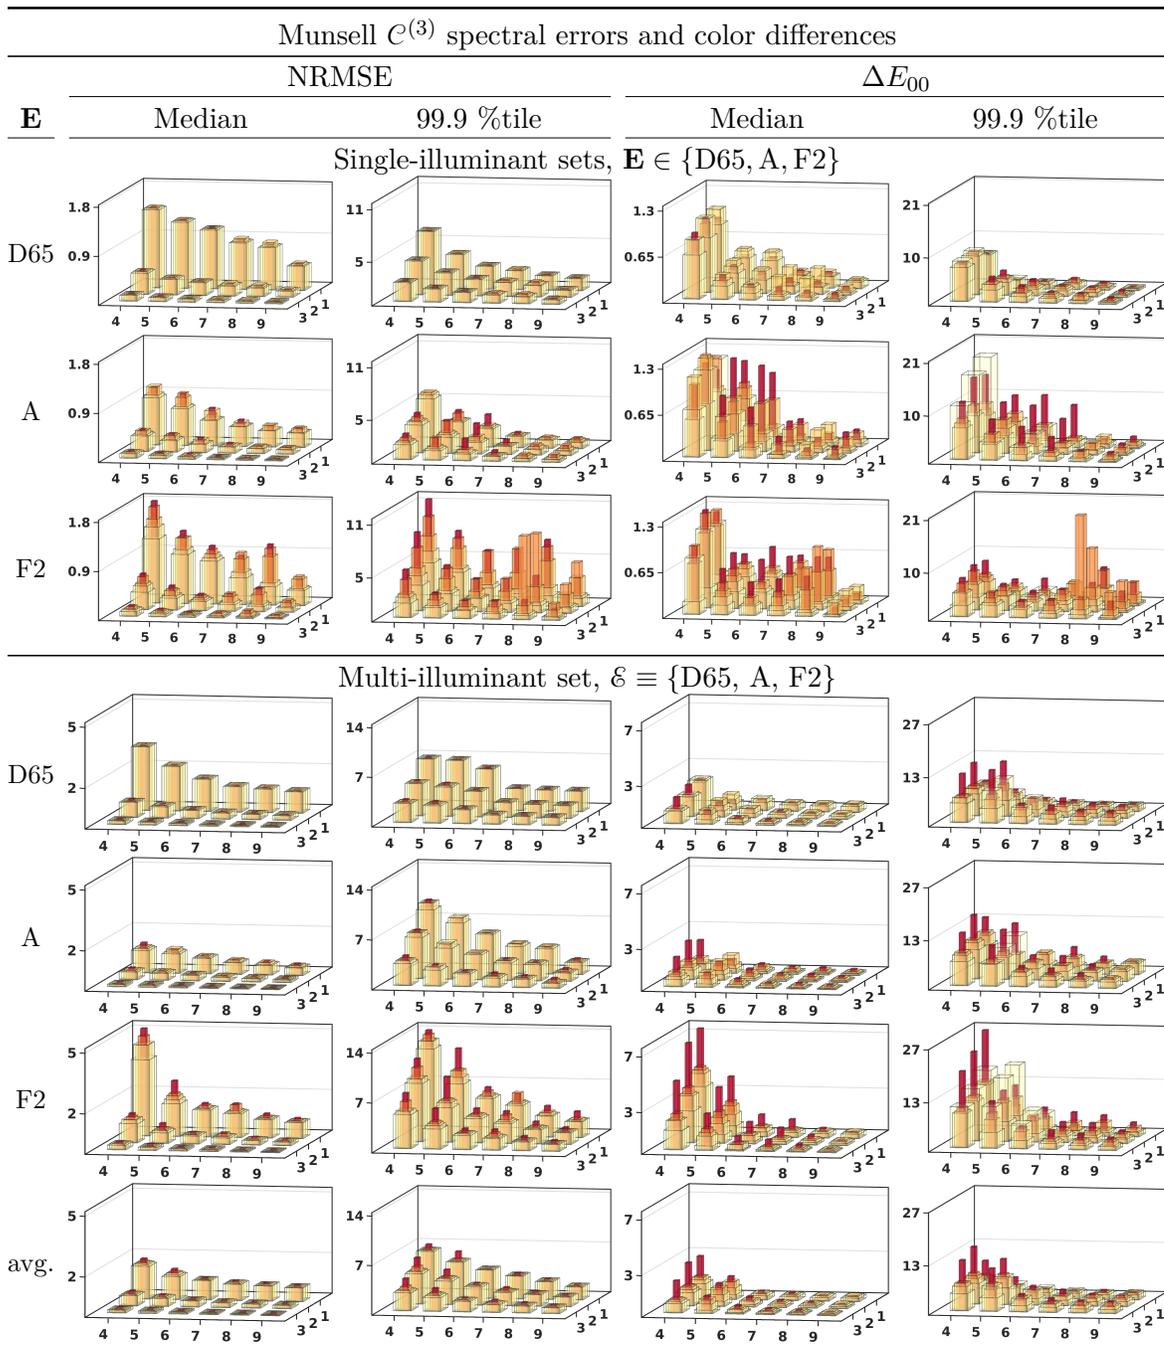


Figure 6.14: For dimensionalities $m \equiv 4, \dots, 9$, and 1st, 2nd and 3rd reflections, the spectral errors and color differences incurred when computing Munsell $\mathcal{C}^{(3)}$ reflections of D65, A, F2 or all three. In the latter case, errors incurred by reflecting the average illuminant SPD are shown in the bottom row. For comparison, the outermost (lightest) bar in each stack of nested bars represents the lower bound error incurred by projections onto $\mathbf{B}^{(2)}$, the first m characteristic vectors of $\mathcal{C}^{(2)}$. The inner four bars represent errors incurred when reflections are computed by multiplying coefficients of sharp bases $\mathbf{Q}^{(i)} \equiv \mathbf{B}^{(2)}\mathbf{T}^{(i)}$, where $\mathbf{T}^{(i)}$ solves, from inner (darker) to outer (lighter) bars, $f^{(5)}$, $f^{(3)}$, $f^{(2)}$ or $f^{(1)}$.

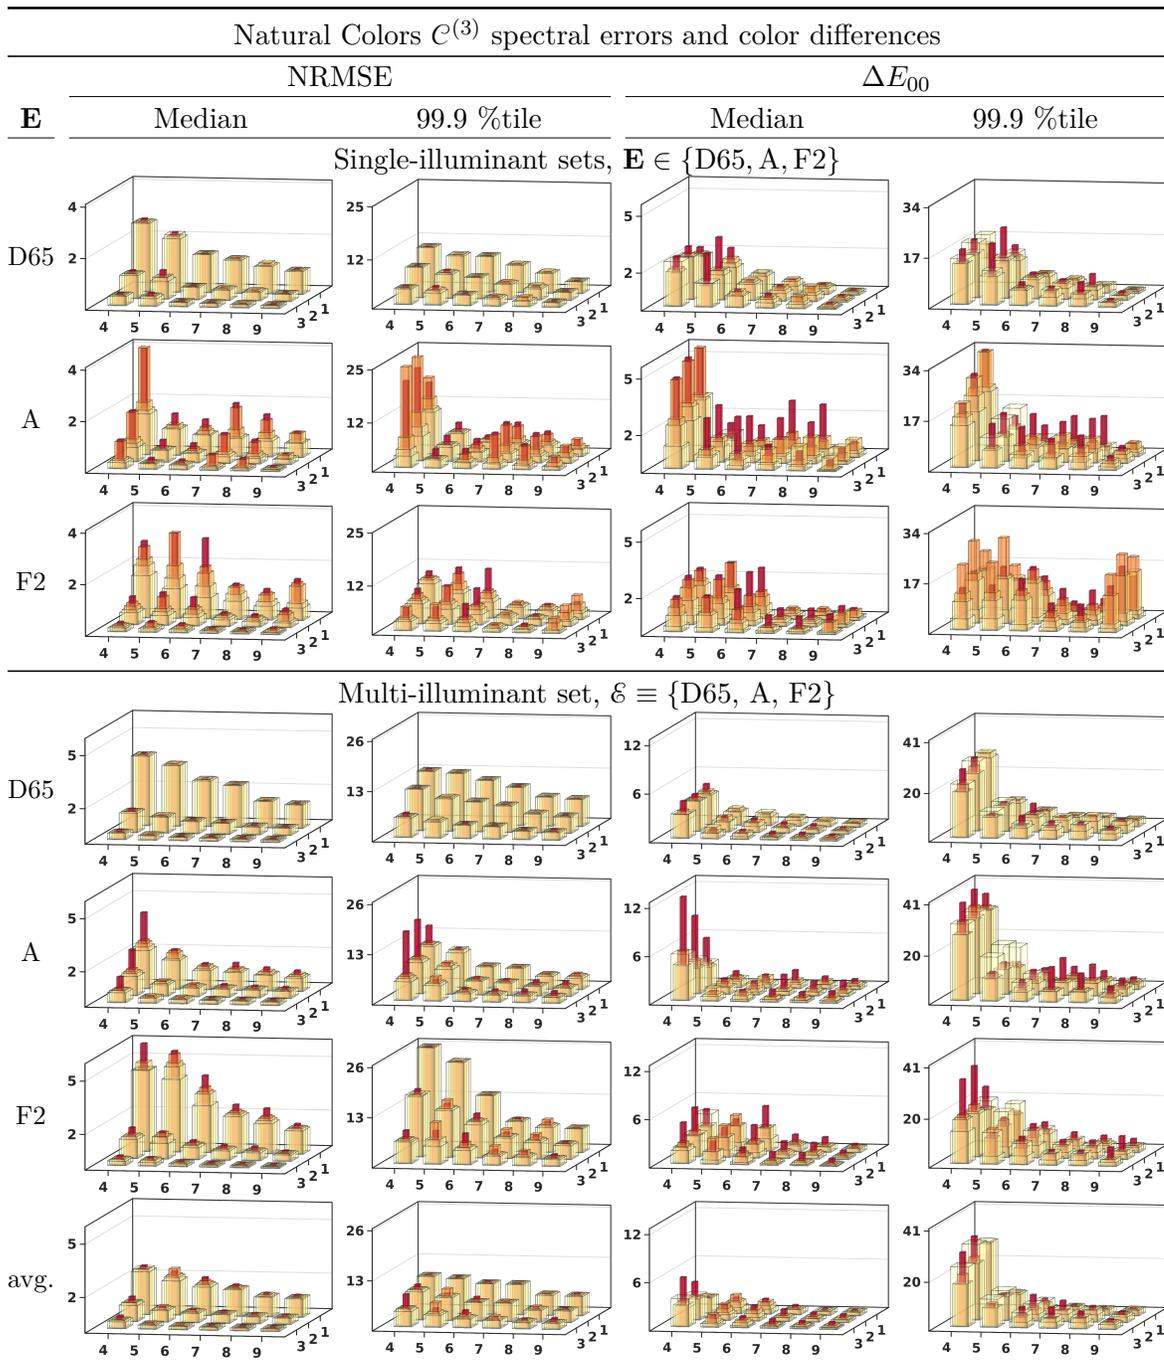


Figure 6.15: For dimensionalities $m \equiv 4, \dots, 9$, and 1st, 2nd and 3rd reflections, the spectral errors and color differences incurred when computing Natural Colors $\mathcal{C}^{(3)}$ reflections of D65, A, F2 or all three. In the latter case, errors incurred by reflecting the average illuminant SPD are shown in the bottom row. For comparison, the outermost (lightest) bar in each stack of nested bars represents the lower bound error incurred by projections onto $\mathbf{B}^{(2)}$, the first m characteristic vectors of $\mathcal{C}^{(2)}$. The inner four bars represent errors incurred when reflections are computed by multiplying coefficients of sharp bases $\mathbf{Q}^{(i)} \equiv \mathbf{B}^{(2)} \mathbf{T}^{(i)}$, where $\mathbf{T}^{(i)}$ solves, from inner (darker) to outer (lighter) bars, $f^{(5)}$, $f^{(3)}$, $f^{(2)}$ or $f^{(1)}$.

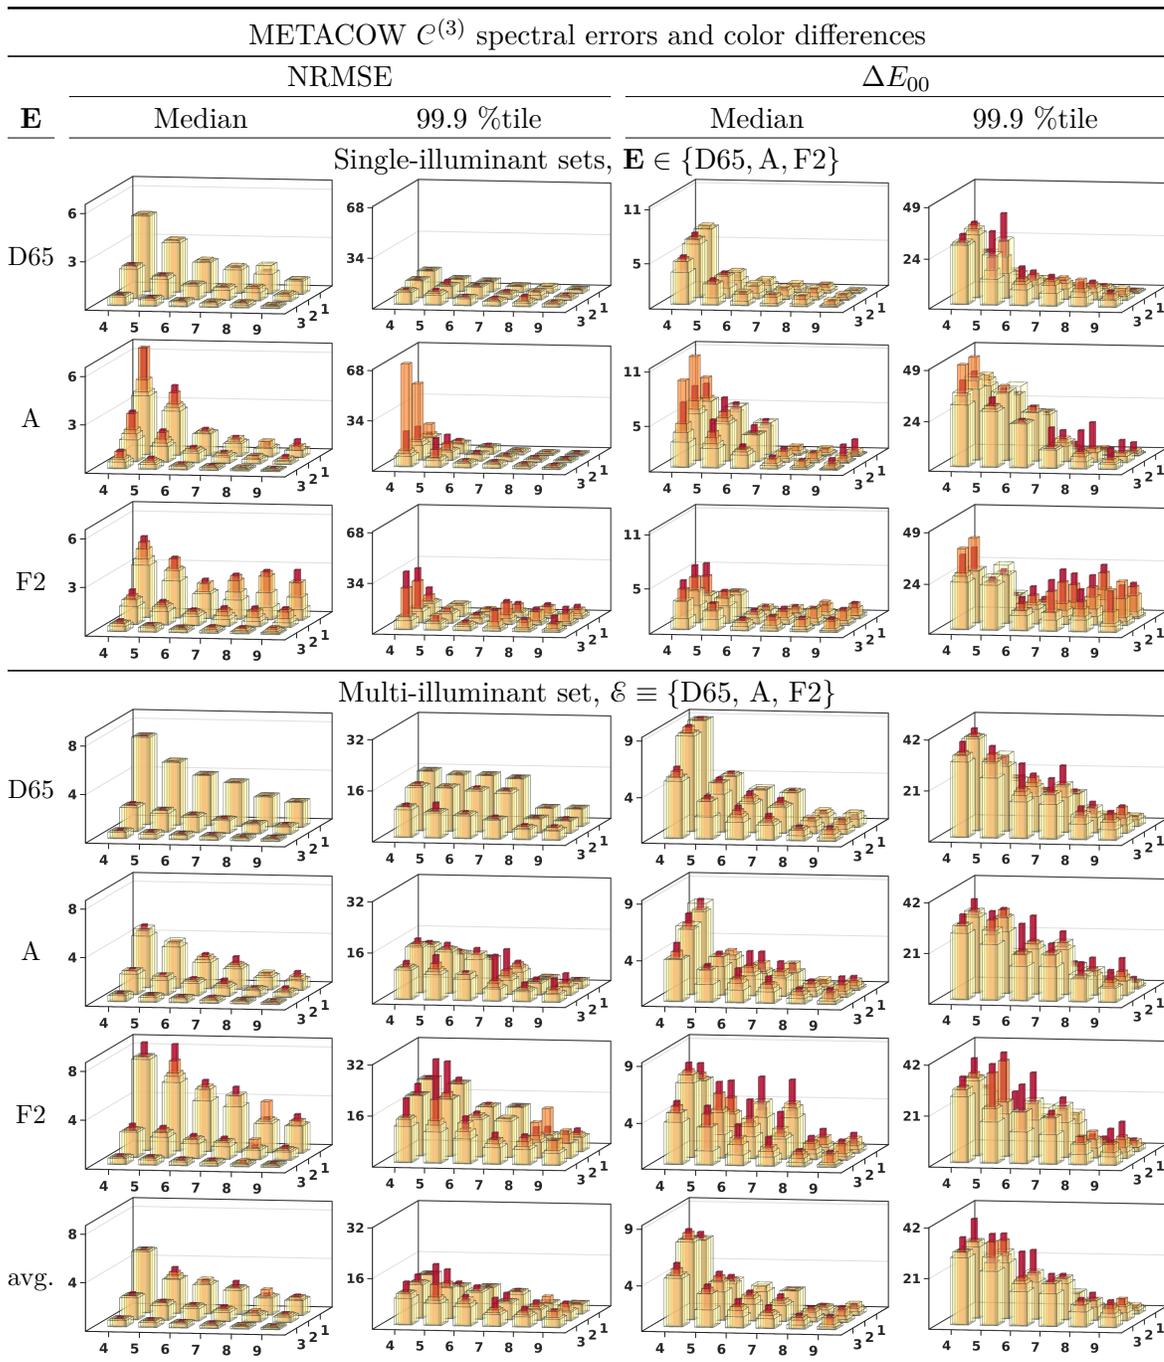


Figure 6.16: For dimensionalities $m \equiv 4, \dots, 9$, and 1st, 2nd and 3rd reflections, the spectral errors and color differences incurred when computing METACOW $\mathcal{C}^{(3)}$ reflections of D65, A, F2 or all three. In the latter case, errors incurred by reflecting the average illuminant SPD are shown in the bottom row. For comparison, the outermost (lightest) bar in each stack of nested bars represents the lower bound error incurred by projections onto $\mathbf{B}^{(2)}$, the first m characteristic vectors of $\mathcal{C}^{(2)}$. The inner four bars represent errors incurred when reflections are computed by multiplying coefficients of sharp bases $\mathbf{Q}^{(i)} \equiv \mathbf{B}^{(2)}\mathbf{T}^{(i)}$, where $\mathbf{T}^{(i)}$ solves, from inner (darker) to outer (lighter) bars, $f^{(5)}$, $f^{(3)}$, $f^{(2)}$ or $f^{(1)}$.

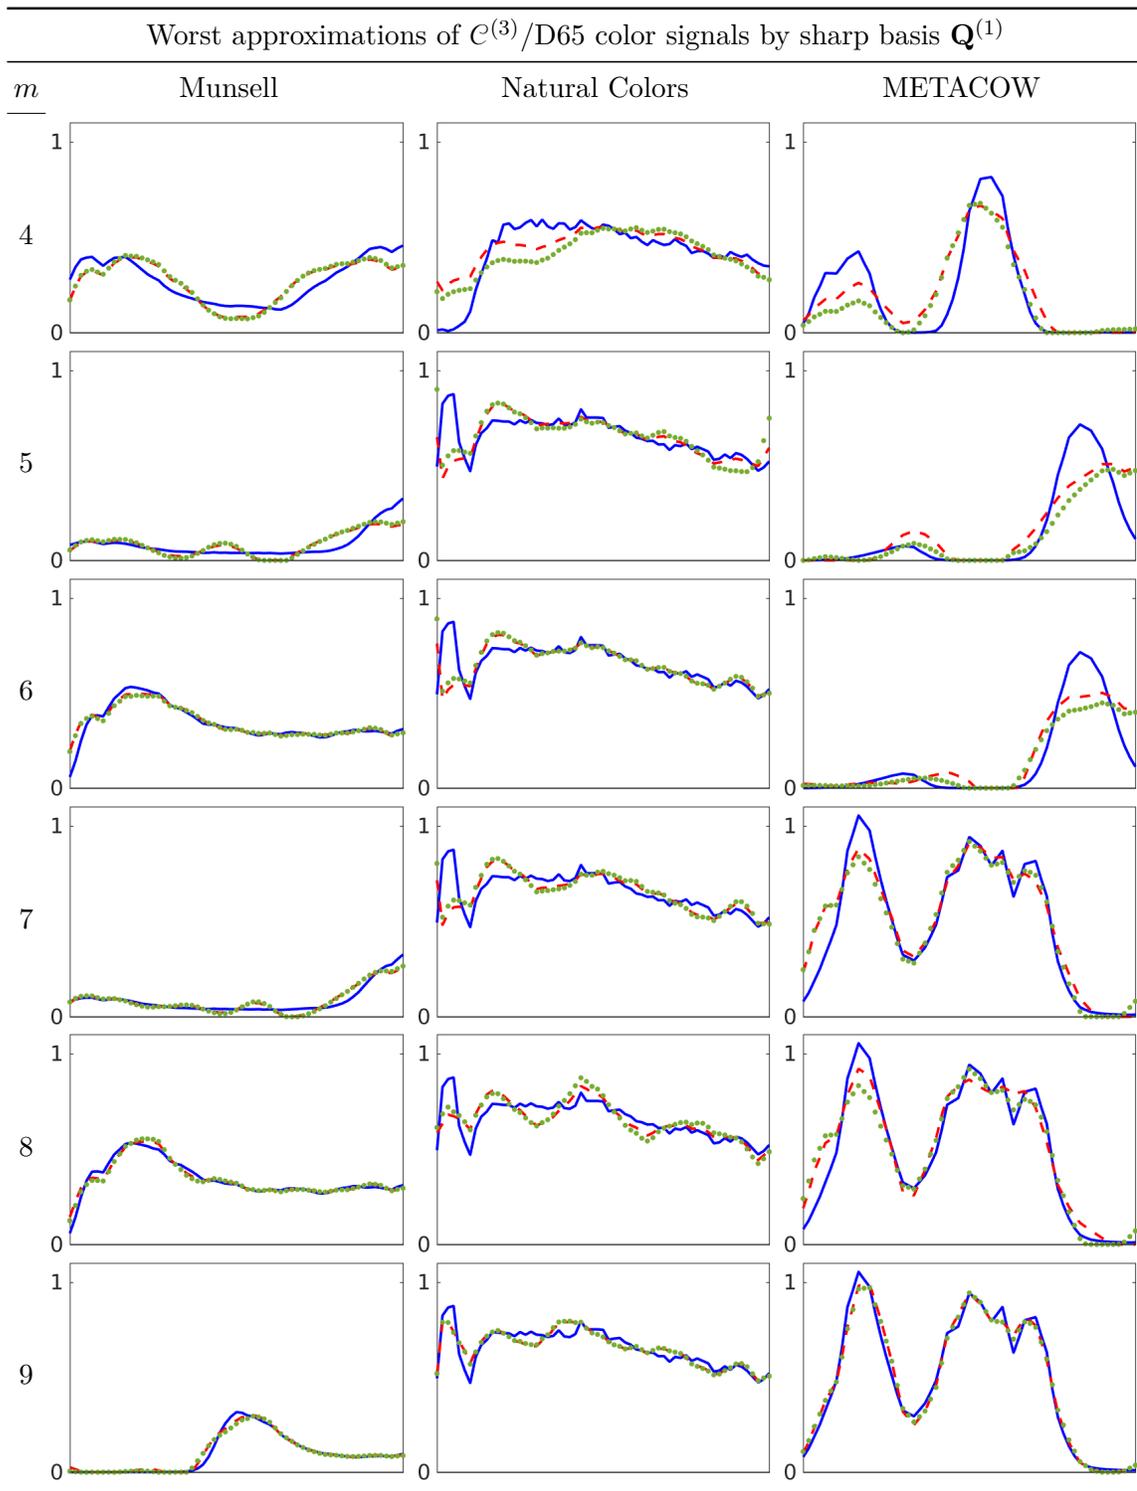


Figure 6.17: From $\mathcal{C}^{(3)}$ sets constructed with illuminant D65, exact color signals are plotted with blue (solid) lines. For dimensionalities $m \equiv 4, \dots, 9$, approximations computed by multiplying coefficients of sharp bases $\mathbf{Q}^{(1)} \equiv \mathbf{B}^{(2)}\mathbf{T}^{(1)}$ are plotted with green (dotted) lines. For comparison, projections onto $\mathbf{B}^{(2)}$ are plotted with red (dashed) lines.

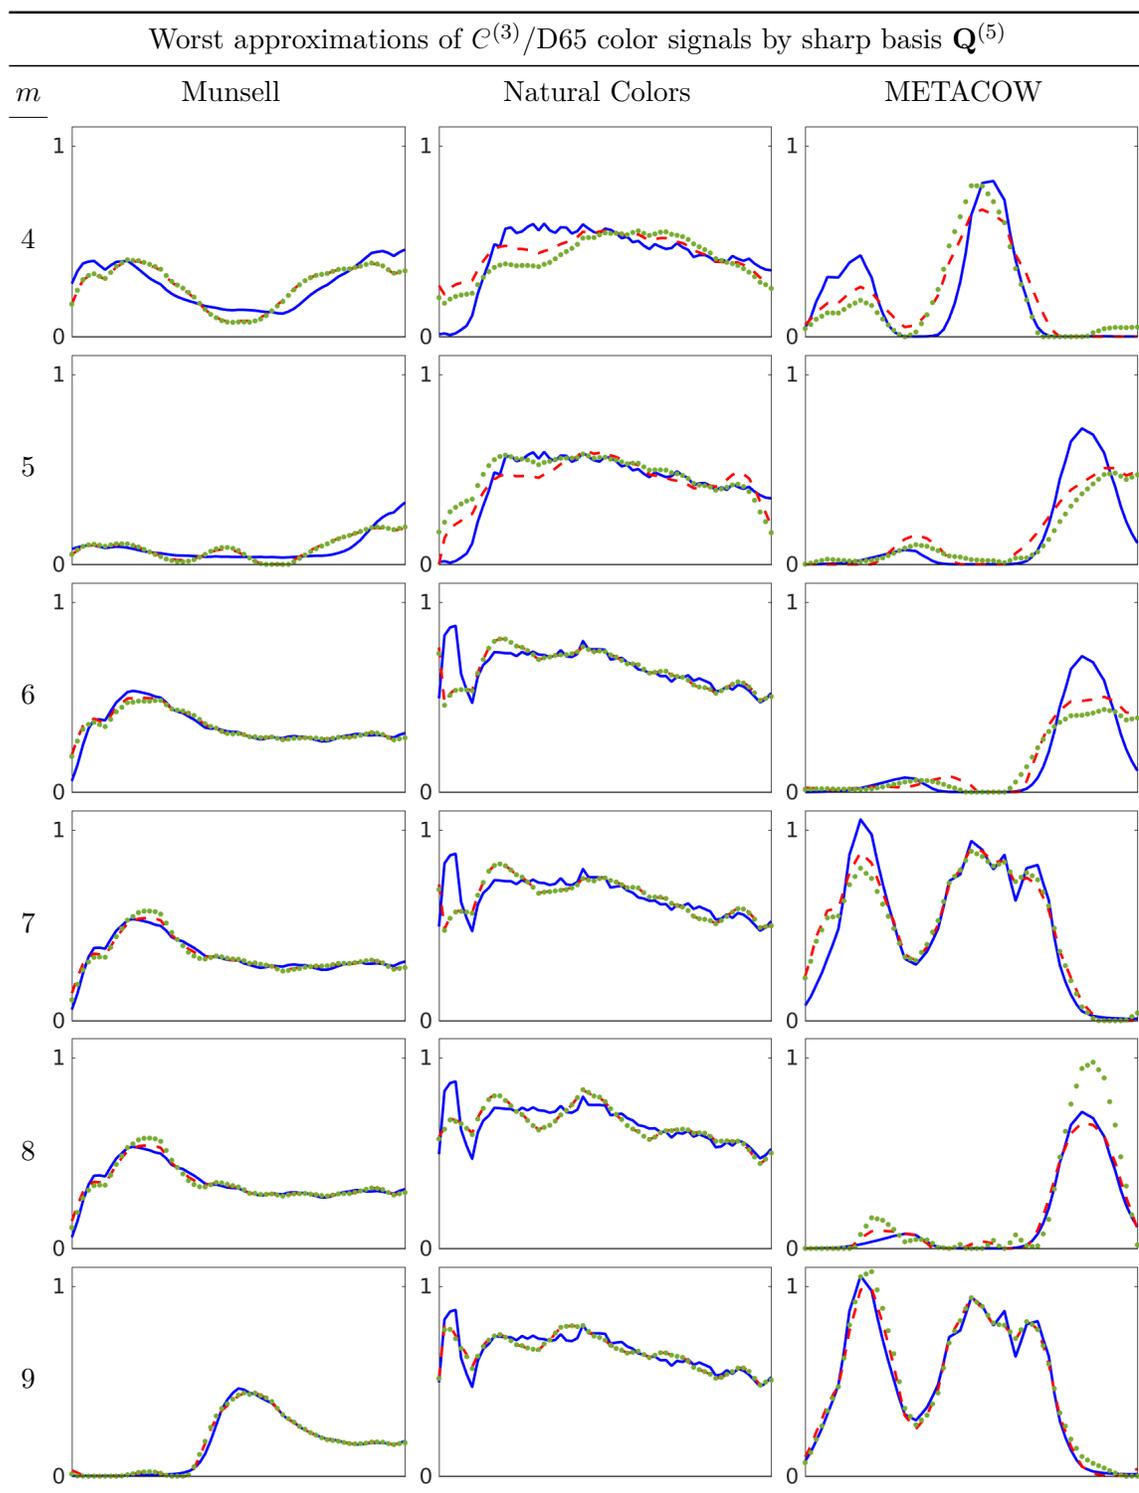


Figure 6.18: From $\mathcal{C}^{(3)}$ sets constructed with illuminant D65, exact color signals are plotted with blue (solid) lines. For dimensionalities $m \equiv 4, \dots, 9$, approximations computed by multiplying coefficients of sharp bases $\mathbf{Q}^{(5)} \equiv \mathbf{B}^{(2)}\mathbf{T}^{(5)}$ are plotted with green (dotted) lines. For comparison, projections onto $\mathbf{B}^{(2)}$ are plotted with red (dashed) lines.

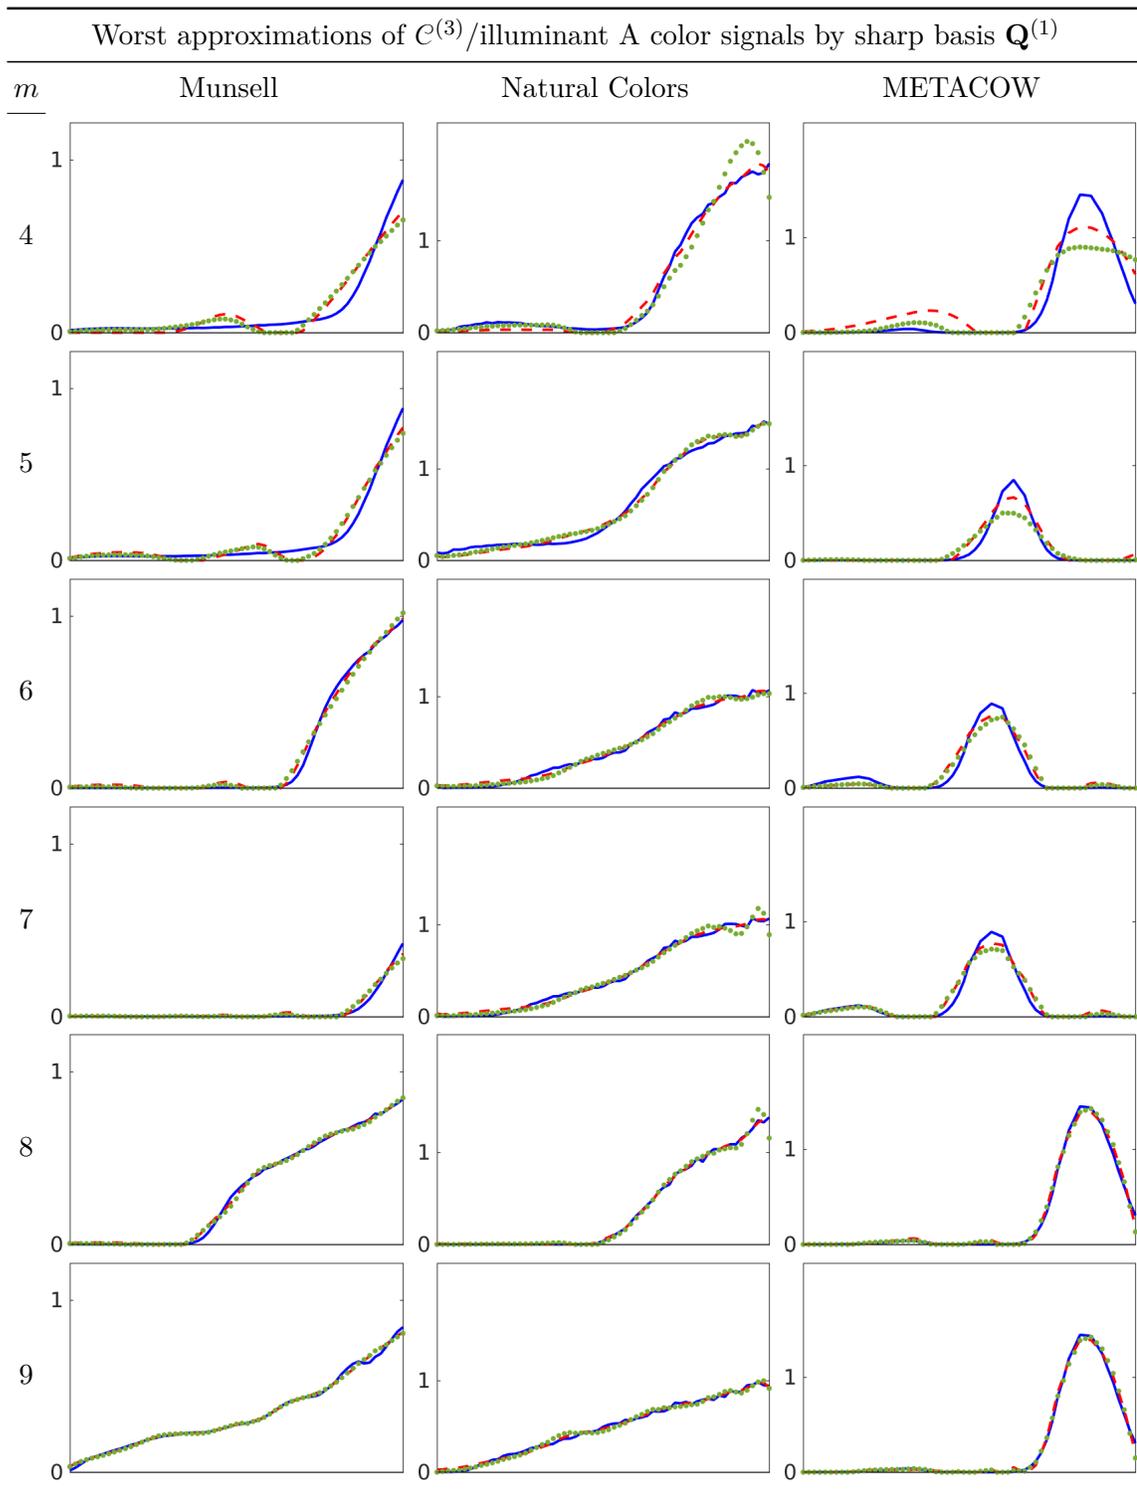


Figure 6.19: From $\mathcal{C}^{(3)}$ sets constructed with illuminant A, exact color signals are plotted with blue (solid) lines. For dimensionalities $m \equiv 4, \dots, 9$, approximations computed by multiplying coefficients of sharp bases $\mathbf{Q}^{(1)} \equiv \mathbf{B}^{(2)}\mathbf{T}^{(1)}$ are plotted with green (dotted) lines. For comparison, projections onto $\mathbf{B}^{(2)}$ are plotted with red (dashed) lines.

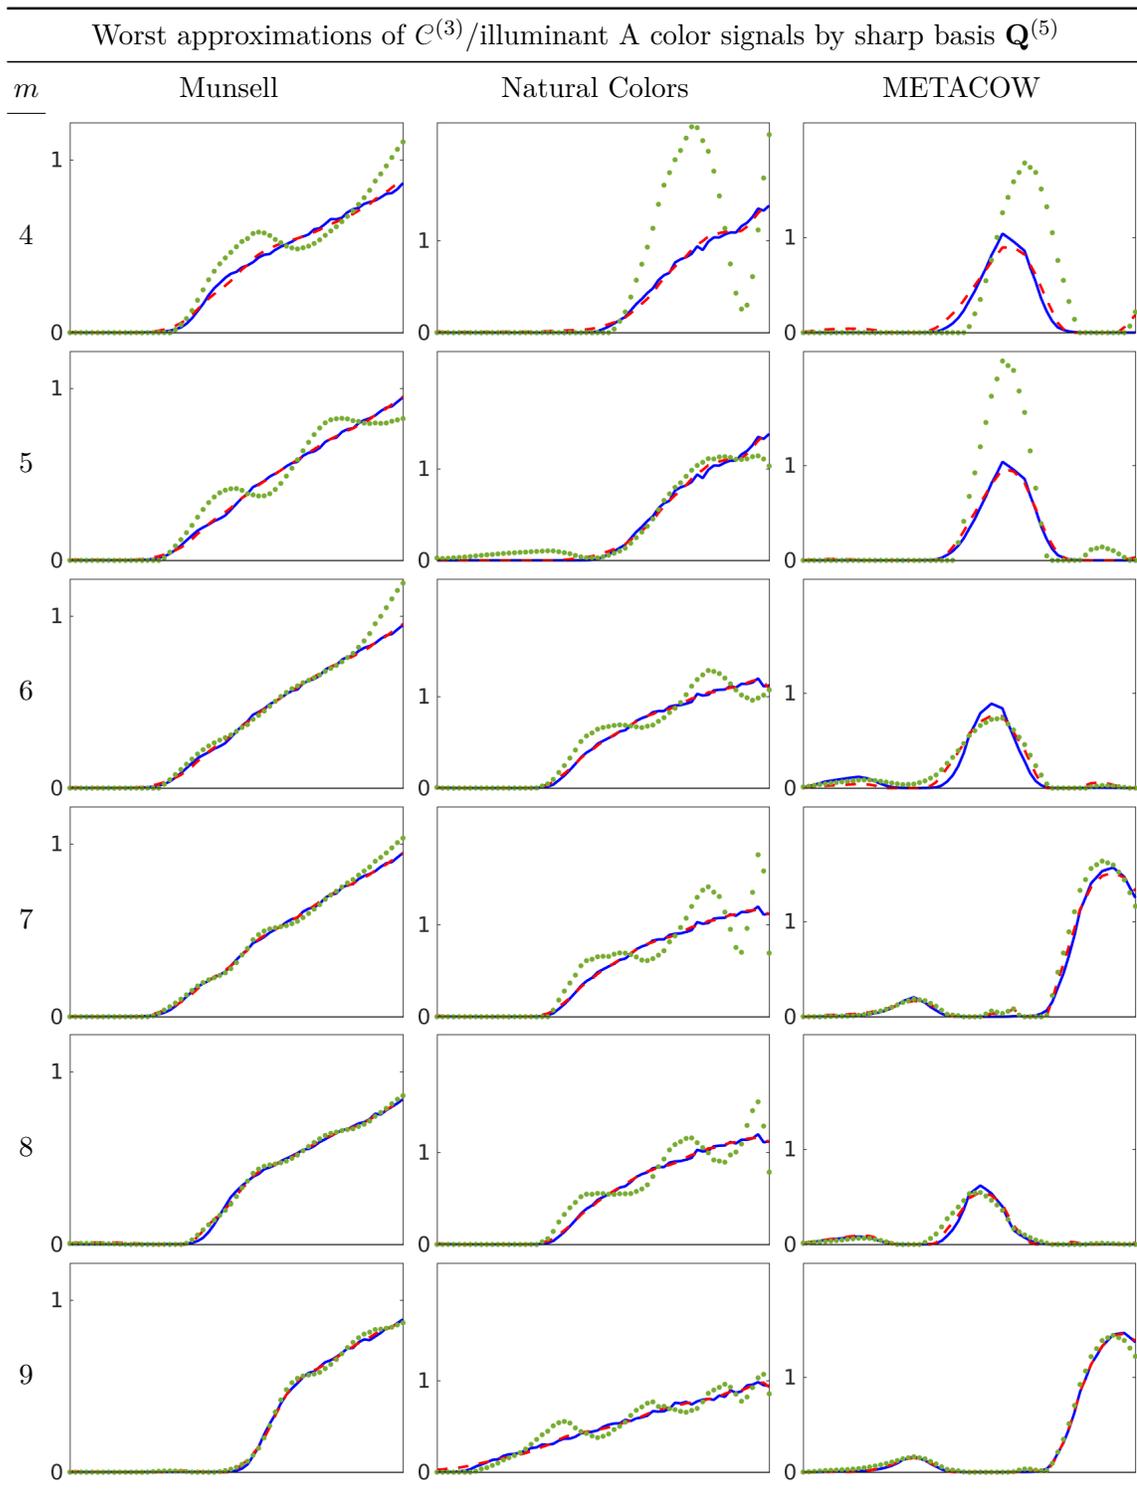


Figure 6.20: From $\mathcal{C}^{(3)}$ sets constructed with illuminant A, exact color signals are plotted with blue (solid) lines. For dimensionalities $m \equiv 4, \dots, 9$, approximations computed by multiplying coefficients of sharp bases $\mathbf{Q}^{(5)} \equiv \mathbf{B}^{(2)}\mathbf{T}^{(5)}$ are plotted with green (dotted) lines. For comparison, projections onto $\mathbf{B}^{(2)}$ are plotted with red (dashed) lines.

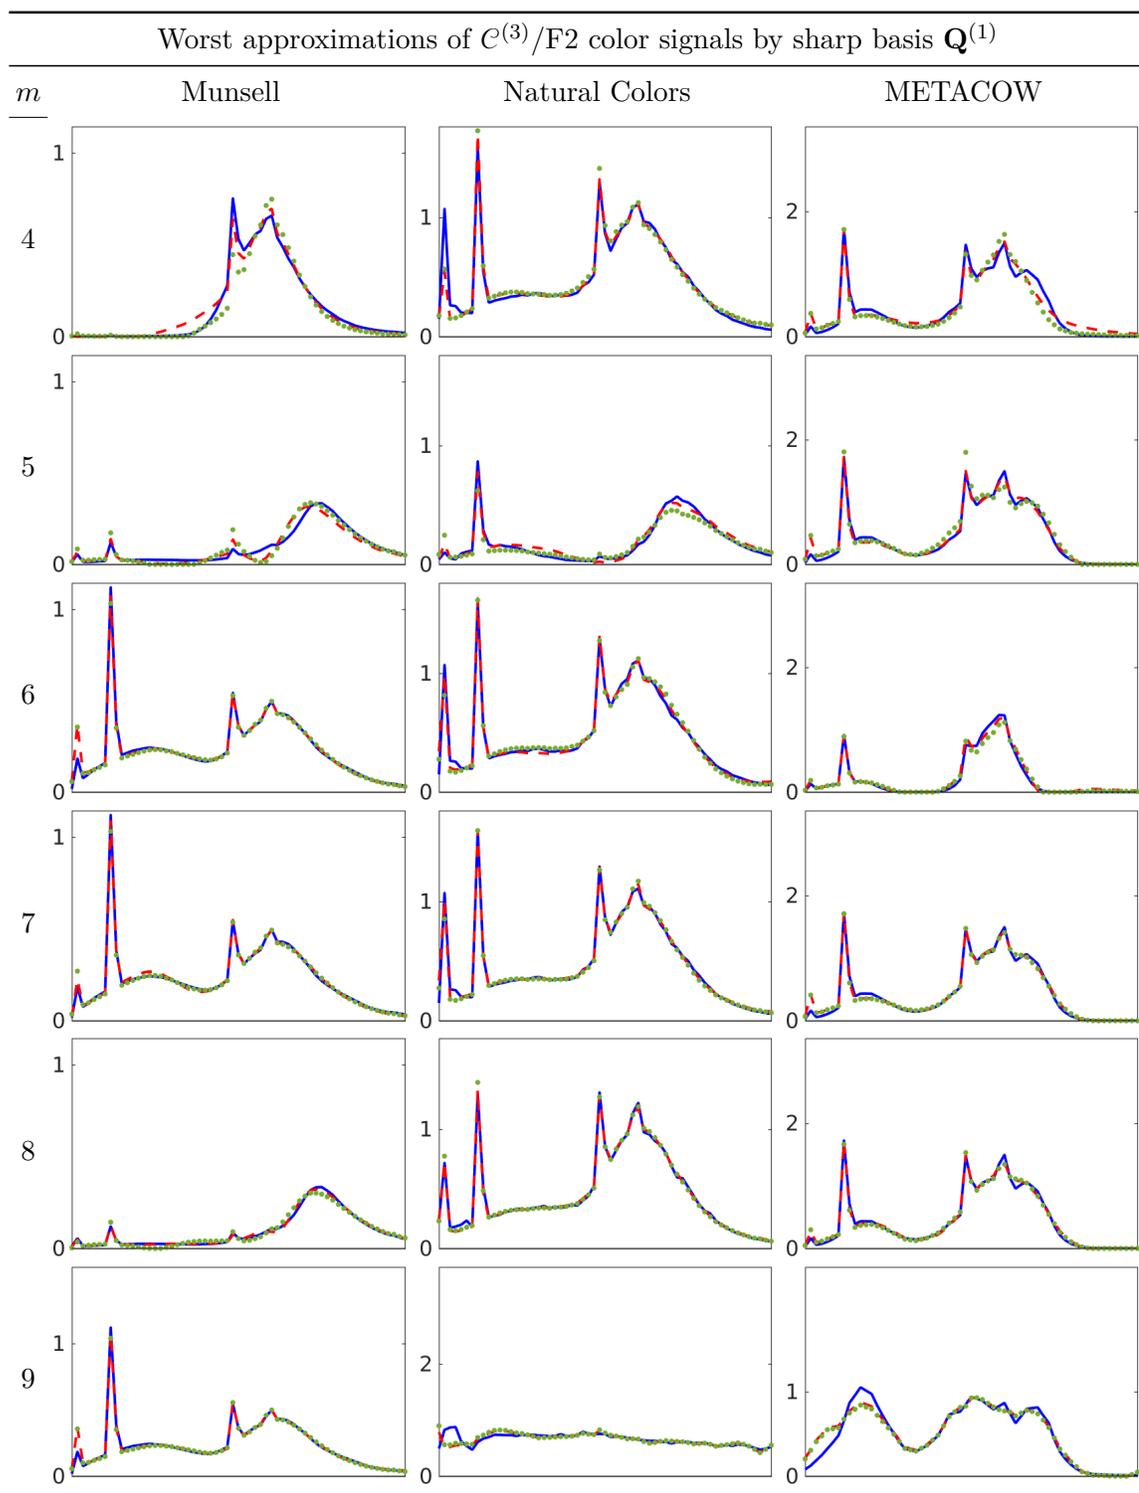


Figure 6.21: From $\mathcal{C}^{(3)}$ sets constructed with illuminant F2, exact color signals are plotted with blue (solid) lines. For dimensionalities $m \equiv 4, \dots, 9$, approximations computed by multiplying coefficients of sharp bases $\mathbf{Q}^{(1)} \equiv \mathbf{B}^{(2)}\mathbf{T}^{(1)}$ are plotted with green (dotted) lines. For comparison, projections onto $\mathbf{B}^{(2)}$ are plotted with red (dashed) lines.

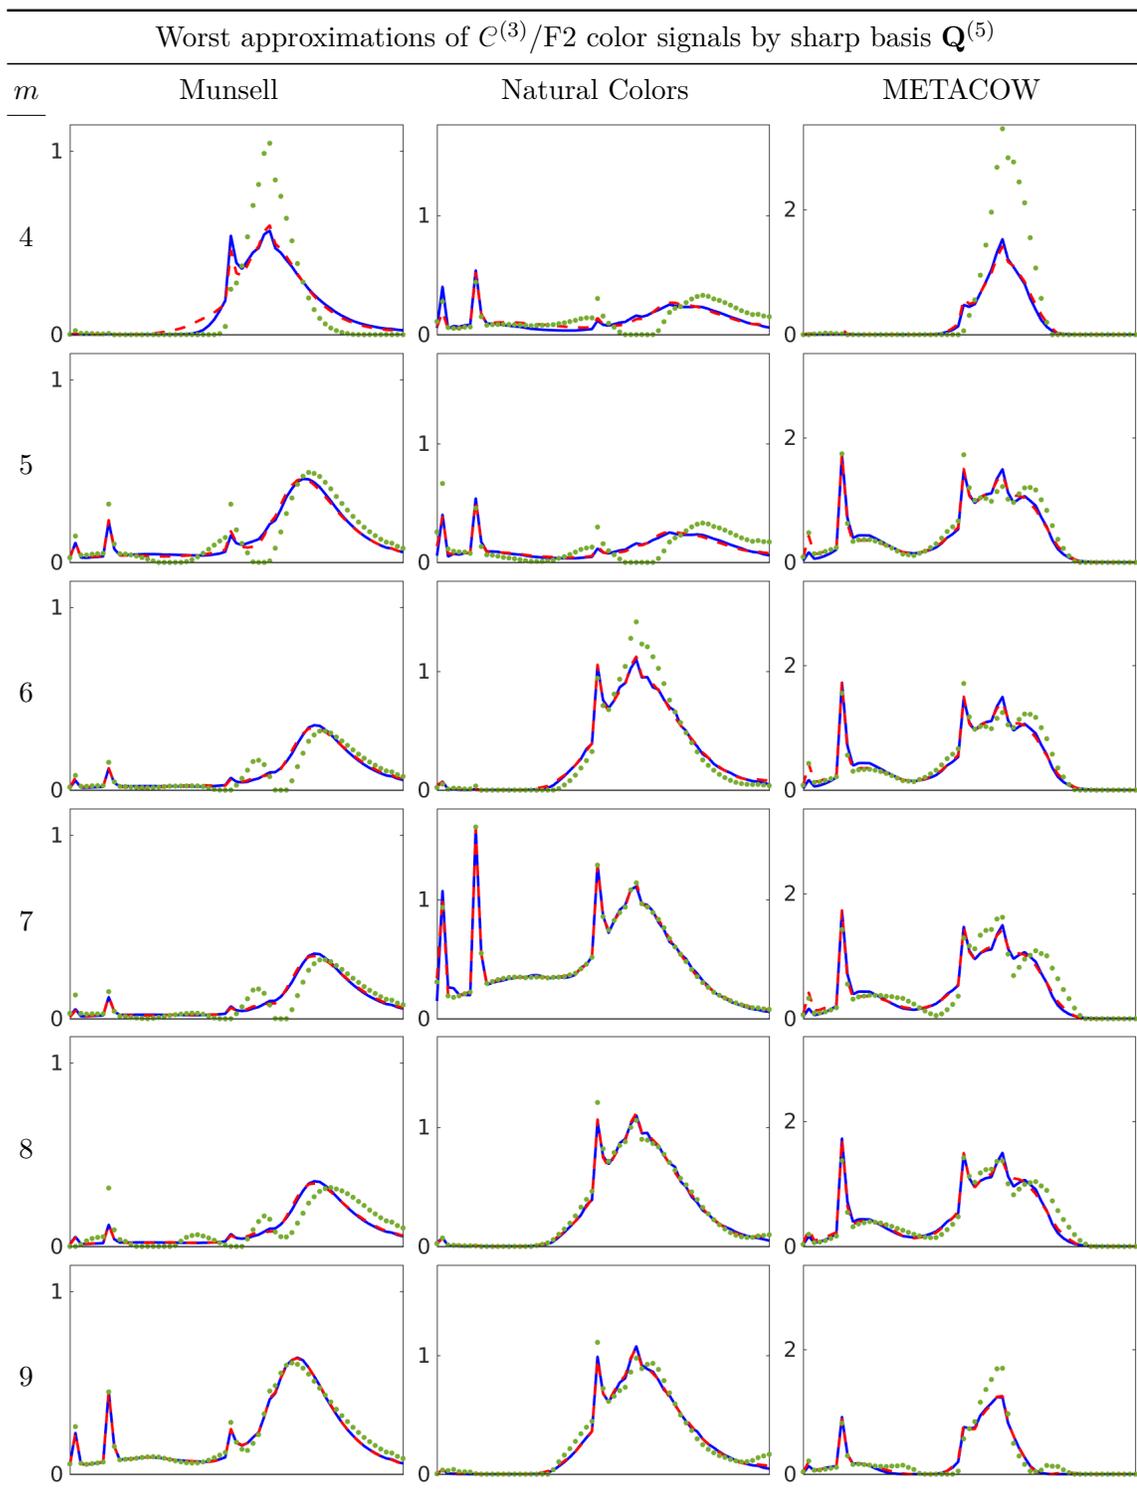


Figure 6.22: From $\mathcal{C}^{(3)}$ sets constructed with illuminant F2, exact color signals are plotted with blue (solid) lines. For dimensionalities $m \equiv 4, \dots, 9$, approximations computed by multiplying coefficients of sharp bases $\mathbf{Q}^{(5)} \equiv \mathbf{B}^{(2)}\mathbf{T}^{(5)}$ are plotted with green (dotted) lines. For comparison, projections onto $\mathbf{B}^{(2)}$ are plotted with red (dashed) lines.

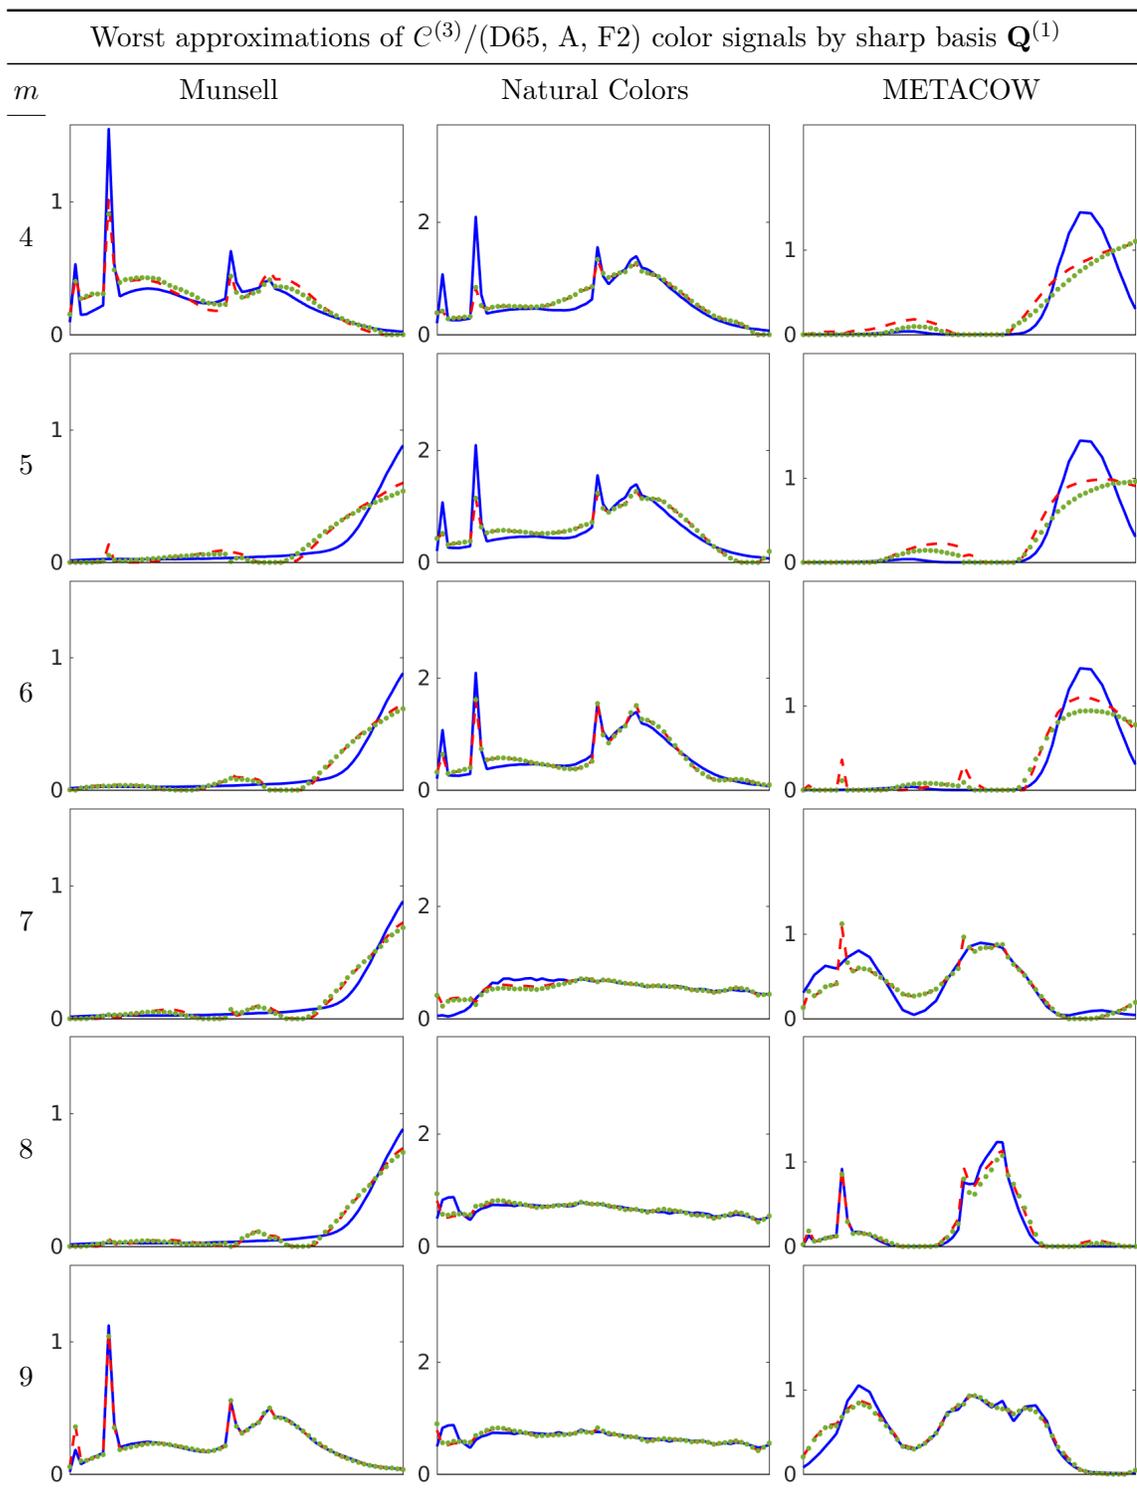


Figure 6.23: From $\mathcal{C}^{(3)}$ sets constructed with illuminants D65, A and F2, exact color signals are plotted with blue (solid) lines. For dimensionalities $m \equiv 4, \dots, 9$, approximations computed by multiplying coefficients of sharp bases $\mathbf{Q}^{(1)} \equiv \mathbf{B}^{(2)}\mathbf{T}^{(1)}$ are plotted with green (dotted) lines. For comparison, projections onto $\mathbf{B}^{(2)}$ are plotted with red (dashed) lines.

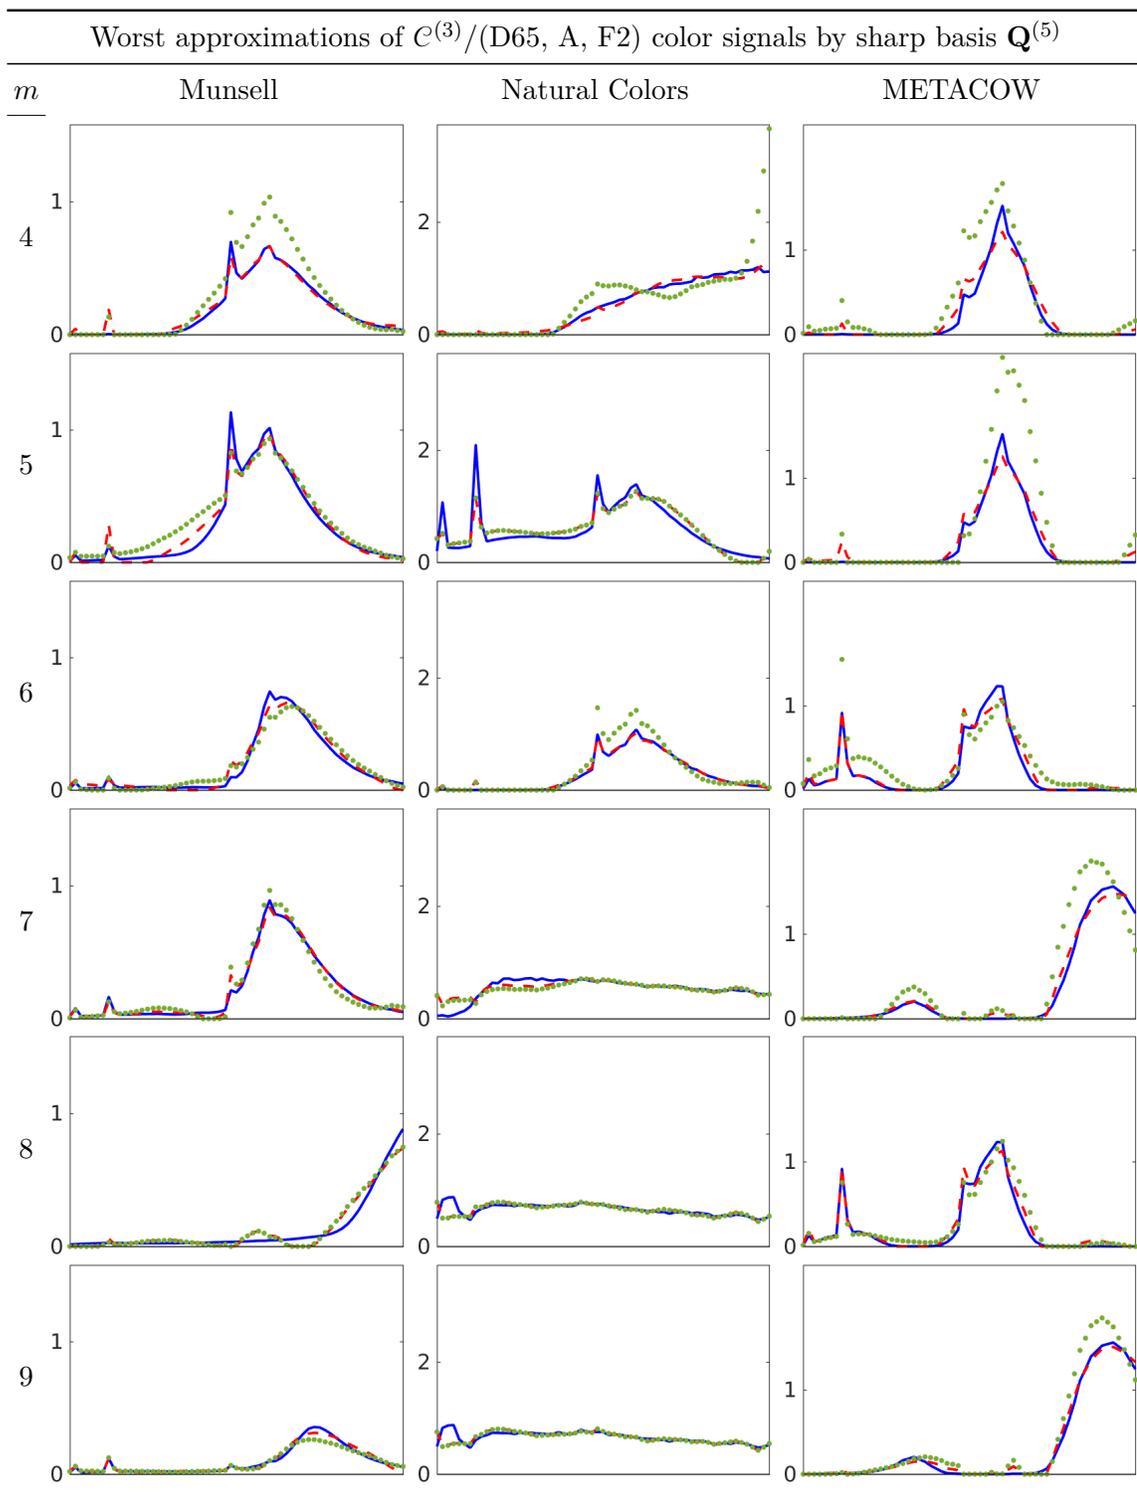


Figure 6.24: From $\mathcal{C}^{(3)}$ sets constructed with illuminants D65, A and F2, exact color signals are plotted with blue (solid) lines. For dimensionalities $m \equiv 4, \dots, 9$, approximations computed by multiplying coefficients of sharp bases $\mathbf{Q}^{(5)} \equiv \mathbf{B}^{(2)}\mathbf{T}^{(5)}$ are plotted with green (dotted) lines. For comparison, projections onto $\mathbf{B}^{(2)}$ are plotted with red (dashed) lines.

6.1.3 Some rendered examples

A demonstration of the scale of CIEDE2000 color differences is provided by the rendered images in Fig. 6.25 (which should be viewed on a monitor, preferably calibrated). This Cornell box scene, illuminated by three different area lights in the ceiling, is rendered by a path tracer, chromatically adapted to D65 and tone-mapped. From the viewer’s perspective, F2 is on the left, D65 is in the center and illuminant A is on the right. The reflectance spectra of the calibration target on the right wall are those of the METACOW image, while those on the left wall were measured from the X-Rite ColorChecker® Digital SG calibration target¹. Using the RGB values of the texture images provided with the vase and tulips model obtained from a commercial 3D model vendor², the flowers’ reflectance spectra are estimated by trilinear interpolation of the spectra available in the Floral Reflectance Database³ [4], augmented with their pointwise squares to enlarge their collective corresponding gamut to include the highly saturated colors.

The large ground truth image was rendered by computing reflections as pointwise products of 61-dimensional spectra. To render the images in the next row, reflections were computed as the products, from left to right, of: 1) RGB triplets of lights and materials; 2) “prefiltered” RGB triplets, as described in [172]; and 3) and 4) coefficients of 4- and 9-dimensional $\mathbf{Q}^{(1)}$ bases. The false color images in the bottom row represent CIEDE2000 color differences, computed as S-CIELAB difference images to approximate the effect of chromatic spatial blurring by the human visual system [182]. Differences in the standard RGB image are easily noticed, while the similar differences in the prefiltered RGB image and the image computed with four-dimensional sharp basis coefficients might be acceptable for some applications. With nine dimensions, differences in the sharp basis image are effectively imperceptible and only very slightly smaller than can be attained with as few as seven dimensions (not shown).

6.2 Real-Time performance

To render a scene using sharp basis coefficients, in real time or otherwise, four one-time pre-processes are required:

¹Spectra provided with ProfileMaker v5.0.10: http://www.xrite.com/service-support/downloads/p/profilemaker_v5_0_10.

²<http://cgaxis.com/product/three-color-tulips/>

³<http://www.reflectance.co.uk/>

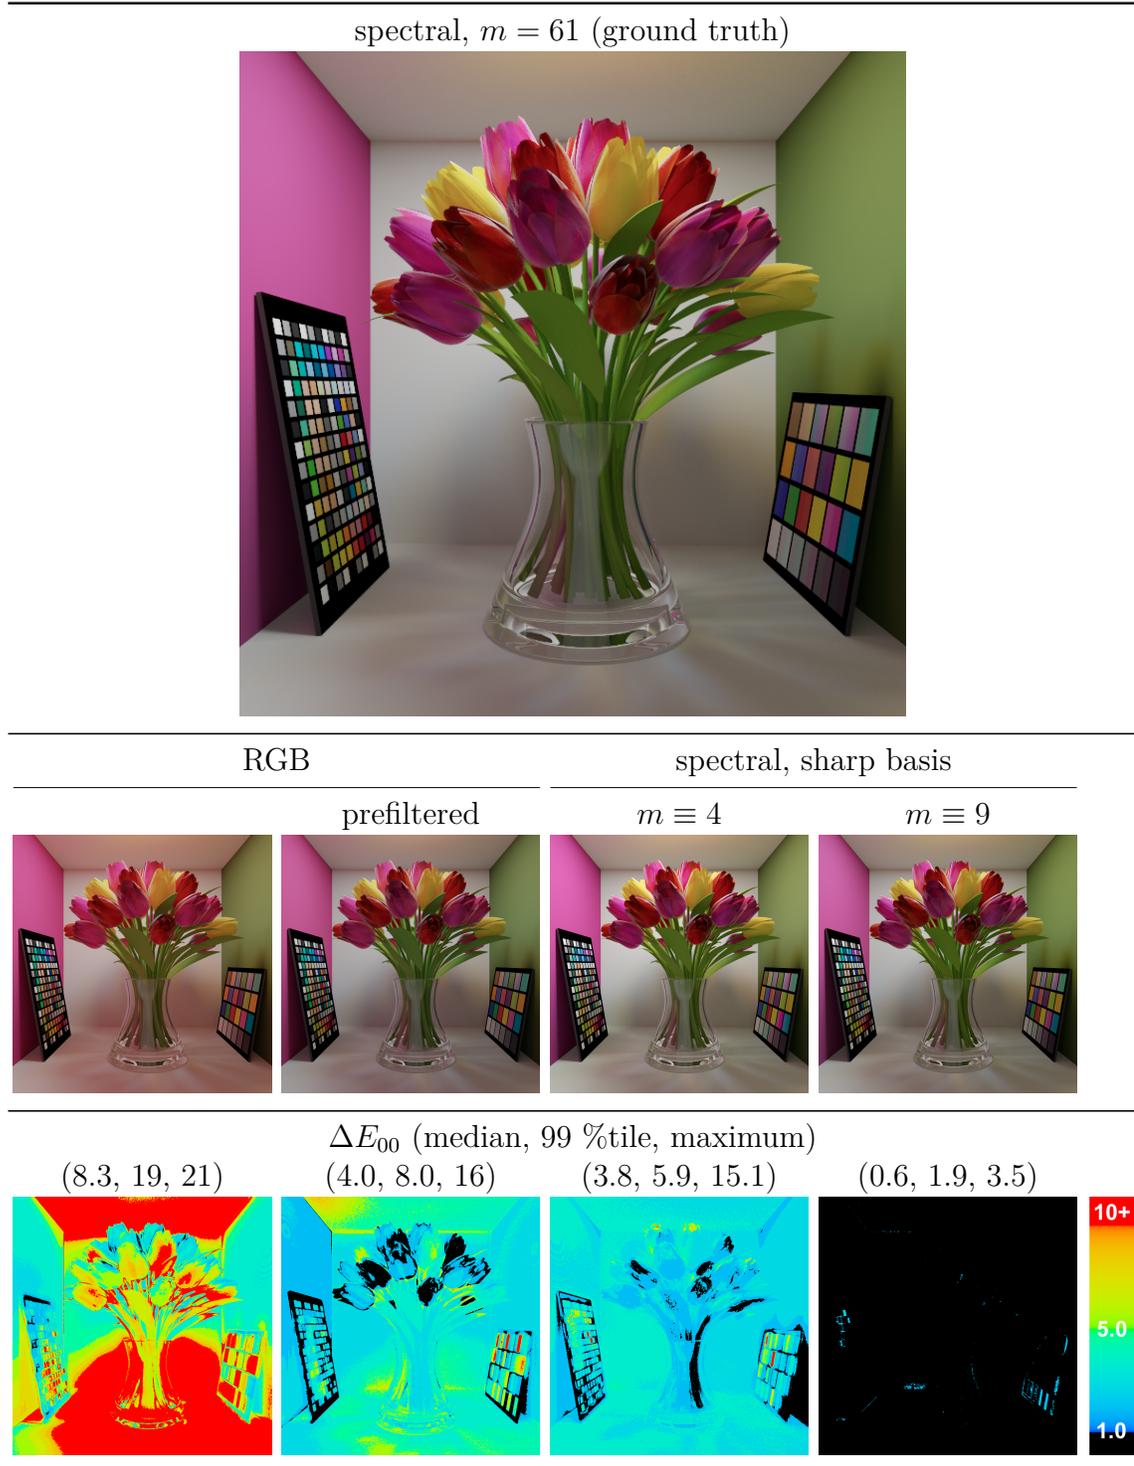


Figure 6.25: Path-traced Cornell box scene, lighted by F2, D65 and illuminant A, chromatically adapted to D65 and tone-mapped. False color S-CIELAB images depict CIEDE2000 color differences between ground truth image on the left and four different methods of color computation. See Section 6.1.3 for details.

1. Construct the matrices \mathcal{S} and \mathcal{E} , filling their columns with, respectively, the reflectance spectra to be rendered and the light source SPDs that will illuminate them. Construct the matrix \mathcal{C} by filling its columns with all $\mathcal{C}^{(k)}$ SPDs, where $k \equiv 1$ when rendering with first reflections only, or $k \equiv 2$ when rendering with indirect illumination.
2. For a given number of dimensions m , chosen for a desired spectral accuracy, construct the orthonormal basis $\mathbf{B}^{(k)}$ from the first m characteristic vectors of \mathcal{C} , obtained from its leading m left singular vectors.
3. Find \mathbf{T} using one of the methods described in Chapter 5.
4. Compute, and save for re-use, the sharp basis coefficients of light source SPDs $\tilde{\mathbf{e}} \equiv \mathbf{T}^{-1}\mathbf{e} \equiv \mathbf{T}^{-1}\mathbf{B}^T\mathcal{E}$ and reflectance spectra $\tilde{\mathcal{J}} \equiv \mathbf{T}^{-1}\mathcal{J}$, where

$$\mathcal{J} \equiv \begin{cases} \mathbf{B}^T \text{diag}(\mathbf{E})^{-1} \mathbf{B} \mathbf{B}^T \text{diag}(\mathbf{E}) \mathcal{S}, & k \equiv 1, q \equiv 1 \\ \mathbf{B}^T \mathcal{S}, & \text{otherwise.} \end{cases}$$

Rendering is then similar to RGB rendering, but using sharp basis coefficients instead of RGB triplets and with a final $3 \times m$ matrix-vector multiplication to compute each pixel's color before display.

In applications in which the illumination is unknown until run time, the above pre-processes can still be carried out in real time, or nearly, depending on the size of \mathcal{S} , by first recursively computing in a pre-process the matrix of unilluminated reflections $\mathcal{S}^{(k)}\mathcal{S}^{(k)T} \equiv \sum_{i=1}^p \text{diag}(\mathcal{S}_i)\mathcal{S}^{(k-1)}\mathcal{S}^{(k-1)T}\text{diag}(\mathcal{S}_i)$, starting with $\mathcal{S}^{(0)}\mathcal{S}^{(0)T} \equiv \mathbf{1}\mathbf{1}^T$. When, during run time, the q light source SPDs \mathcal{E} are known (where typically q is small), the basis $\mathbf{B}^{(k)}$ can be found in less than a millisecond by computing the $n \times n$ symmetric matrix $\mathcal{C}^{(k)}\mathcal{C}^{(k)T} \equiv \sum_{i=1}^q \text{diag}(\mathcal{E}_i)\mathcal{S}^{(k)}\mathcal{S}^{(k)T}\text{diag}(\mathcal{E}_i)$ and taking its leading m singular vectors. After quickly finding \mathbf{T} and computing \mathbf{Q} and \mathbf{Q}^+ , the bulk of this runtime process is consumed by computing the sharp coefficients $\tilde{\mathcal{J}}$, which, for $m \equiv 4$ to 9 and a million reflectance spectra, can be done on a CPU in approximately 90 to 140 milliseconds. On a GPU, coefficients would likely be computed in parallel in a tenth of this time or less.

In real-time rendering, the relative cost of shading fragments with sharp basis coefficients is $\lceil \frac{m}{\alpha} \rceil$, where α depends on the use of the alpha channel. In scenes with only opaque materials, $\alpha \equiv 4$, because an RGBA vector can store four coefficients;

Table 6.1: Relative rendering costs

	<i>m</i>						
RGB	4	5	6	7	8	9	61
1.0	1.2	1.2	1.2	1.5	1.5	1.6	20

Relative costs of rendering of a 1024×768 textured full-screen quad, for dimensionality $m \equiv 61$ by multiplying spectra componentwise, and for $m \equiv 4, \dots, 9$ by multiplying sharp basis coefficients. The alpha channel was reserved for transparency (i.e., $\alpha \equiv 3$).

otherwise, $\alpha \equiv 3$. The cost of computing RGB coordinates is the cost of reading coefficients from the render texture(s) plus the cost of a $3 \times m$ matrix-vector multiplication at each pixel. Because the memory read will dominate, the total cost of rendering a frame should approximate a step function of m with a width of 3 or 4.

Real-time rendering performance can be limited by CPU computation or by vertex or fragment shading on the GPU. A conservative estimate of the real cost of rendering with sharp basis coefficients would therefore be the cost incurred by a scene that: 1) requires minimal CPU computation; 2) renders as little geometry as possible; 3) has a simple lighting model; and 4) reads the reflectance properties of every fragment from GPU memory. Table 6.1 shows the costs, relative to RGB rendering, of the rendering of a 1024×768 textured, full-screen quad with only ambient lighting. Although the costs shown here range from 20% to 60% above the cost of RGB rendering, the actual cost in a typical application would be significantly lower or even insignificant.

Chapter 7

Conclusion

7.1 Summary

The usual goal of realism in computer graphics is best defined by its opposite. For an image to be plausible, it need only not be implausible. More specifically, since *realism* in this context means *photorealism*, the question “Does this look real?” means “Does this look like a photograph?” which leaves as much latitude in meeting this goal as there is in capturing a scene with a camera. *Predictive rendering*, on the other hand, ask a different question [102, 122, 176, 73, 127]: “Given these initial conditions (i.e., this arrangement of objects in space, these optical properties of materials, this illumination, etc.), how will this scene appear?” Implicit in this question is a particular observer, which could be an abstract, empirically constructed, human *standard observer* [56], a non-human animal [110, 111, 94, 8], or a particular camera or type of film [87, 58, 86]. Also implicit is a context. Is the predicted appearance that of a photograph? Is it printed on paper, projected or displayed on a particular computer monitor? In what surroundings? Or is such an image meant somehow *to represent* the observer’s actual visual sensation, and, if so, in what sense?

Clearly, *appearance* is difficult to define and the problem of its prediction is easily ill-posed. Nevertheless, any prediction must start with models of light and the optical properties of a scene’s materials, for which RGB triplets are inadequate. The next simplest model are spectra, and their interaction can be simply modeled as the pointwise products of n -dimensional vectors when simulating the transport of light through a scene. Implementation of this straightforward approach mimics RGB rendering, but in n dimensions instead of three, with a final $3 \times n$ matrix-vector multiplication at

each pixel to yield an image ready for various transformations before display. This approach is also useful, of course, in contexts in which the distribution of light itself in a scene, not its appearance, is the subject of interest. This is the case, for example, when modeling the transport of invisible infrared or ultraviolet light. The methods described in the preceding chapters allow this transport to be simulated in real time applications by representing spectra in $m \ll n$ dimensions. They improve on the spectral sharpening method of Drew and Finlayson [49] by finding change of basis matrices \mathbf{T} that yield more accurate approximations. They also find them more quickly.

Typical users will be more motivated by the $O(m)$ cost of computing reflections than by the $p \cdot (n-m)/n$ savings in data storage, the cost of which continues to decline. These users, however, will be limited to those needing to render images quickly, either in real time or in large batches, since computing reflections in n dimensions is both simpler and more accurate. Compared to the almost ubiquitous presence in the world today of imagery generated using the RGB heuristic, applications using spectral rendering to predict color appearance are rare, notwithstanding the examples cited in the introduction. The popularization of virtual reality, however, could change this, as accurate modeling of color appearance could make environments more immersive. Still, another hurdle remains, as spectral data are scarce and may remain so until hyperspectral imaging is as inexpensive and convenient as digital photography. In the meantime, applications will likely be specialized and limited to those needing to render relatively few spectra, obtained from public sources, or those critical enough to justify the high cost of data acquisition, such as sophisticated simulators used to train pilots or surgeons. For any use, however, and for a given computational cost, the methods presented in this work likely offer the best possible accuracy.

7.2 Future Work

Mentioned in Section 3.1.1, dispersive materials having indices of refraction that vary by wavelength are typically rendered using some version of ray tracing, either by brute force, with a large number of rays representing single-wavelength photons, or using a method to select, stochastically or otherwise, refracted rays at dispersive interfaces. A low-dimensional, linear model of spectra is not well suited for rendering highly dispersive materials, which in the extreme case would require accurate approximation of a delta function. However, a composite model, like Sun's [158], could. Such

a model would incorporate two kinds of rays: those attached to relatively smooth spectra, which can be represented as sharp basis coefficients, and single-wavelength photons described by wavelength and power. Alternatively, non-smooth rays could transport multiple wavelengths chosen stochastically, as in Wilkie’s hero wavelength method [174]. Implementation of such a model would determine if the overhead of maintaining two classes of rays is offset by gained efficiency and flexibility.

Other wavelength-dependent effects deserving investigation are interference and fluorescence. An example of the former is iridescence caused by a thin film, which reflects incident light that has interfered with itself coherently after a small phase shift induced after reflection by two closely spaced surfaces, amplifying the power at some wavelengths while diminishing others, depending on the angle of view and the distance between the surfaces. To render iridescence, the effect of this interference can be precomputed and stored with sufficient resolution in a three-dimensional table of reflectance spectra represented as sharp basis coefficients and indexed by distance and view angle. The question, however, is whether they can be accurately approximated in a small number of dimensions, since they can be much sharper than normal reflectance spectra, which is why iridescent colors are usually saturated. This question could be easily answered, without implementation, by measuring some number of mathematically generated spectra of typical iridescent materials. When rendering normal materials, the error incurred by multiplying sharp basis coefficients, which accumulates and propagates on secondary reflections, will be sufficiently offset by its diminished power. This may not be the case for iridescent materials, whose reflectance spectra can exceed unity, which is why iridescent colors are also often bright. This question too can be answered by the same method, without implementation. The rendering of fluorescent materials, which reflect light at wavelengths different from those of the incident light, poses a different problem. The question is whether such a reflectance, which is inherently a matrix-vector operation, can be sufficiently diagonalized to be represented as coefficients of an m -dimensional sharp basis or, if necessary, as an $m \times m$ matrix. In either case, fluorescent materials may require more dimensions to approximate, since their colors, like iridescent colors, tend to be saturated.

On the theoretical side, a fundamental question is unanswered by the analysis presented in Chapter 5. The lower bound on the total approximation error, $f^{(1)} \geq \sum_{i=1}^n \|(\mathbf{B}\mathbf{B}^T - \mathbf{I}) \text{diag}(\hat{\mathbf{U}}_i^{(s)})\mathcal{C}\|_F^2$, is clear and understood. An upper bound can be

obtained in relation to $f^{(5)}$, our measure of the sharpness of \mathbf{Q} :

$$\begin{aligned}
f^{(1)} &\leq \|\mathbf{c}\|_F^2 \|\Delta\|_F^2 \\
&\equiv \|\mathbf{c}\|_F^2 \sum_{k=1}^n \|\mathbf{B}^T \text{diag}(\widehat{\mathbf{U}}_k^{(\mathcal{S})}) \mathbf{B} - \mathbf{T} \text{diag}(\widetilde{\mathbf{u}}_k^{(\mathcal{S})}) \mathbf{T}^{-1}\|_F^2 \\
&\leq \|\mathbf{c}\|_F^2 \|\kappa_F(\mathbf{T})\| \sum_{k=1}^n \|\mathbf{Q}^+ \text{diag}(\widehat{\mathbf{U}}_k^{(\mathcal{S})}) \mathbf{Q} - \text{diag}(\widetilde{\mathbf{u}}_k^{(\mathcal{S})})\|_F^2 \\
&\leq \|\mathbf{c}\|_F^2 \|\kappa_F(\mathbf{T})\| \|\mathcal{S}\|_F^2 [\|\mathbf{Q}^{+T} \circ \mathbf{Q} - \mathbf{Q}^{+T}\|_F^2 + \sum_{i \neq j} \|(\mathbf{Q}^{+T})_i \circ \mathbf{Q}_j\|_F^2] \\
&\approx \|\mathbf{c}\|_F^2 \|\kappa_F(\mathbf{T})\| \|\mathcal{S}\|_F^2 [\|\mathbf{Q}^{+T} \circ \mathbf{Q} - \mathbf{Q}^{+T}\|_F^2 + \sum_{i \neq j} \frac{\|\mathbf{Q}_i\|_F^2 \|\mathbf{Q}_j\|_F^2}{(\mathbf{Q}_{i,1})^2 (\mathbf{Q}_{j,1})^2} \|\mathbf{Q}_i \circ \mathbf{Q}_j\|_F^2] \quad (7.2) \\
&\approx \|\mathbf{c}\|_F^2 \|\kappa_F(\mathbf{T})\| \|\mathcal{S}\|_F^2 [\|\mathbf{Q}^{+T} \circ \mathbf{Q} - \mathbf{Q}^{+T}\|_F^2 + \sum_{i \neq j} \|\mathbf{Q}_i \circ \mathbf{Q}_j\|_F^2] \\
&= \|\mathbf{c}\|_F^2 \|\kappa_F(\mathbf{T})\| \|\mathcal{S}\|_F^2 [\|\mathbf{Q}^{+T} \circ \mathbf{Q} - \mathbf{Q}^{+T}\|_F^2 + f^{(5)}] \xrightarrow{f^{(5)} \rightarrow 0} 0,
\end{aligned}$$

where $\kappa_F(\mathbf{T})$ is the Frobenius condition number of \mathbf{T} , which for sufficiently sharp \mathbf{Q} is near its lower bound of 1, since $\mathbf{T}^{-1} \approx \text{diag}(\widehat{\mathbf{Q}}^T \mathbf{1})^{-1} \widehat{\mathbf{T}}^T$, where $\widehat{\mathbf{Q}} \equiv \mathbf{B} \widehat{\mathbf{T}}$ and $\widehat{\mathbf{T}}^T \widehat{\mathbf{T}} = \mathbf{I}$. Also for sufficiently sharp \mathbf{Q} , this bound is tight, since, in the limit, where $\mathbf{Q}^{+T} \circ \mathbf{Q} = \mathbf{0}$, we have $\mathbf{Q}^+ = \text{diag}(\widehat{\mathbf{Q}}^T \mathbf{1})^{-1} \widehat{\mathbf{Q}}^T$. In this case, (7.2) is an equality, and we have $\mathbf{Q}^{+T} \circ \mathbf{Q} \approx \mathbf{Q}^{+T}$ (see Section 5.3) and $f^{(5)} \approx f^{(1)} \approx 0$.

This is a formal way of stating the obvious, which is that the error incurred when multiplying sharp basis coefficients is related to the sharpness of \mathbf{Q} and vanishes in the limit where its columns are disjoint. The unanswered question is: What is lowest upper bound on $f^{(5)}$, as a function of \mathbf{B} ? In other words, how much can a given orthonormal basis be sharpened and which aspect(s) of its structure does this depend on? Since a random basis cannot be sharpened, it might likely depend in some way on the basis vectors' quasi-periodicity.

Consider, for example, the plots in Fig. 7.1. Four-dimensional orthonormal bases \mathbf{B} are shown in the left column, starting at the top with a discrete cosine sequence and continuing to the bottom with the leading four characteristic vectors of Munsell $\mathcal{C}^{(1)}$ color signals formed with illuminants $E \equiv \mathbf{1}$, D65, A and F2. In the right columns are plots of the corresponding maximally sharp $\mathbf{Q}^{(5)}$ superimposed on the columns/rows of the projection matrix $\mathbf{B}\mathbf{B}^T$ plotted in gray. Notably, the basis for the E color signals, which are the Munsell reflectance spectra, closely resembles the discrete cosine sequence, which is expected of the eigenvectors of the correlation

matrix of highly correlated data. Since the near-constant D65 SPD has little effect on the reflectance spectra, the D65 basis bears a similar resemblance. This is, of course, less true of the illuminant A basis and even less true of the F2 basis, although they both retain a noticeable quasi-periodicity. Also noticeable, and most important here, is that all maximally sharp $\mathbf{Q}^{(5)}$ basis vectors can be chosen, or very nearly, from among the columns (rows) of the corresponding projection matrix, each of which resembles, from top to bottom by greater to lesser degrees, D_4 , the fourth order Dirichlet kernel (or partial sum of the first four discrete cosines), with phases shifted and amplitudes modulated by the corresponding rows (columns). It is empirically evident that a bound on $f^{(5)}$, which in turn bounds $f^{(1)}$, should be conceivable as some function of a mathematically defined feature, or features, of \mathbf{B} . If such a bound can be found, and if the bound is tight, then, given a set of reflectance spectra and some number of illuminant SPDs, the overall error incurred when rendering with sharp basis coefficients can be approximated in advance and the method's accuracy guaranteed.

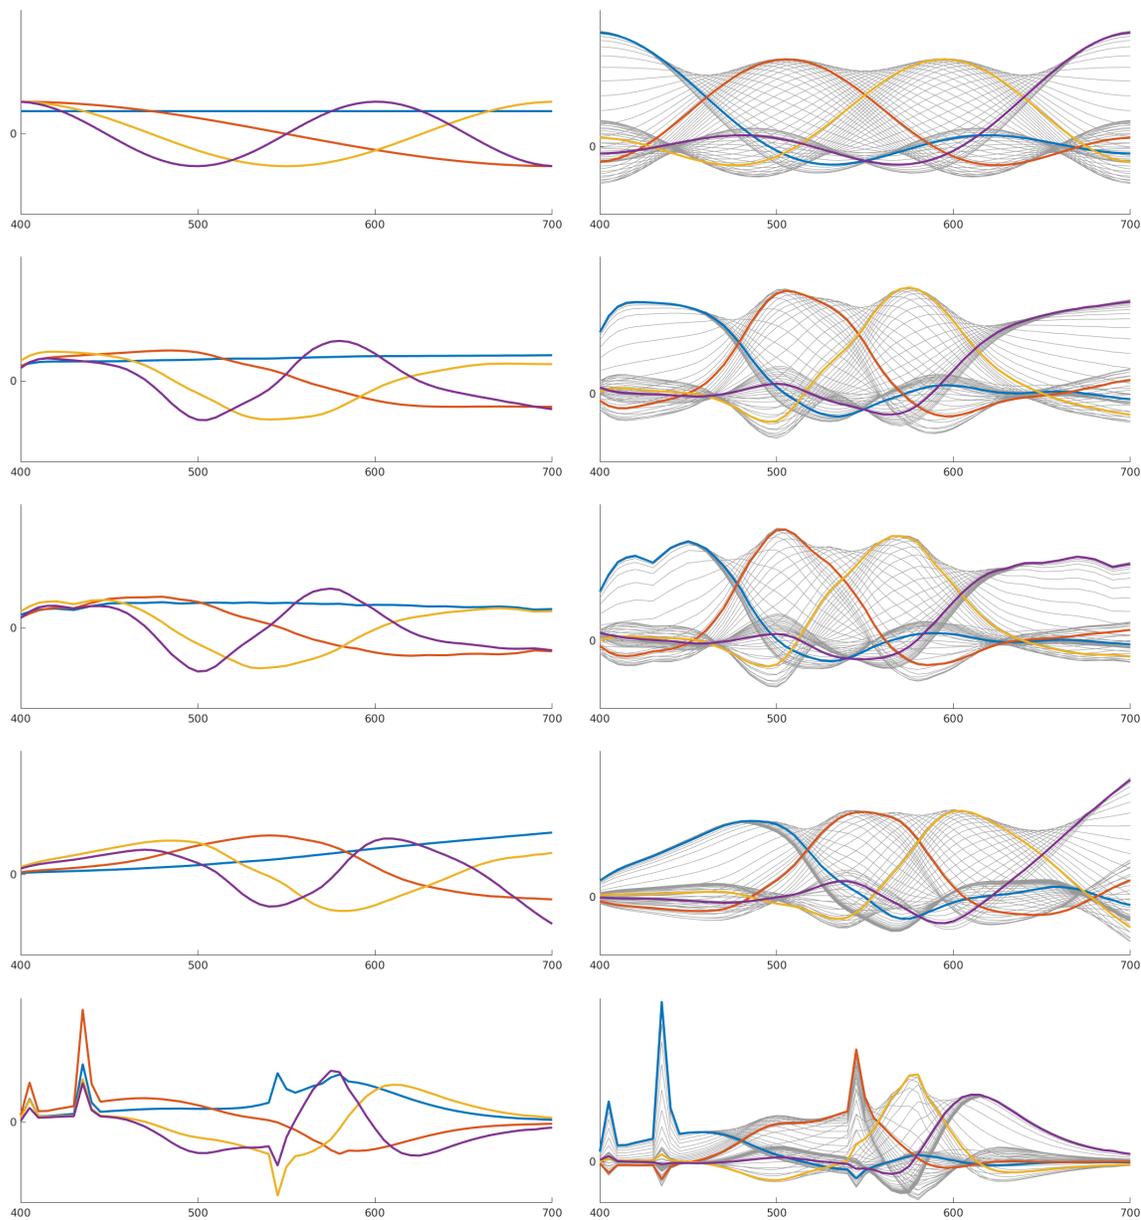


Figure 7.1: In left column are orthonormal bases \mathbf{B} formed with, from top to bottom, a discrete cosine sequence and the leading m characteristic vectors of Munsell $\mathcal{C}^{(1)}$ color signals formed with illuminants E, D65, A and F2. In the right column are plots of the corresponding maximally sharp $\mathbf{Q}^{(5)}$ superimposed on the gray columns/rows of projection matrix $\mathbf{B}\mathbf{B}^T$. All vectors are normalized to unit length.

Appendix

The following figures show detailed distributions of $\mathbf{Q}^{(1)}$ and $\mathbf{Q}^{(5)}$ spectral errors and color differences, as well as, for comparison, those incurred by projection onto \mathbf{B} . Figs. A1-A12 show distributions for single-illuminant $\mathcal{C}^{(1)}$ SPDs. Distributions for all $\mathcal{C}^{(3)}$ SPDs are shown in Figs. A13-A36.

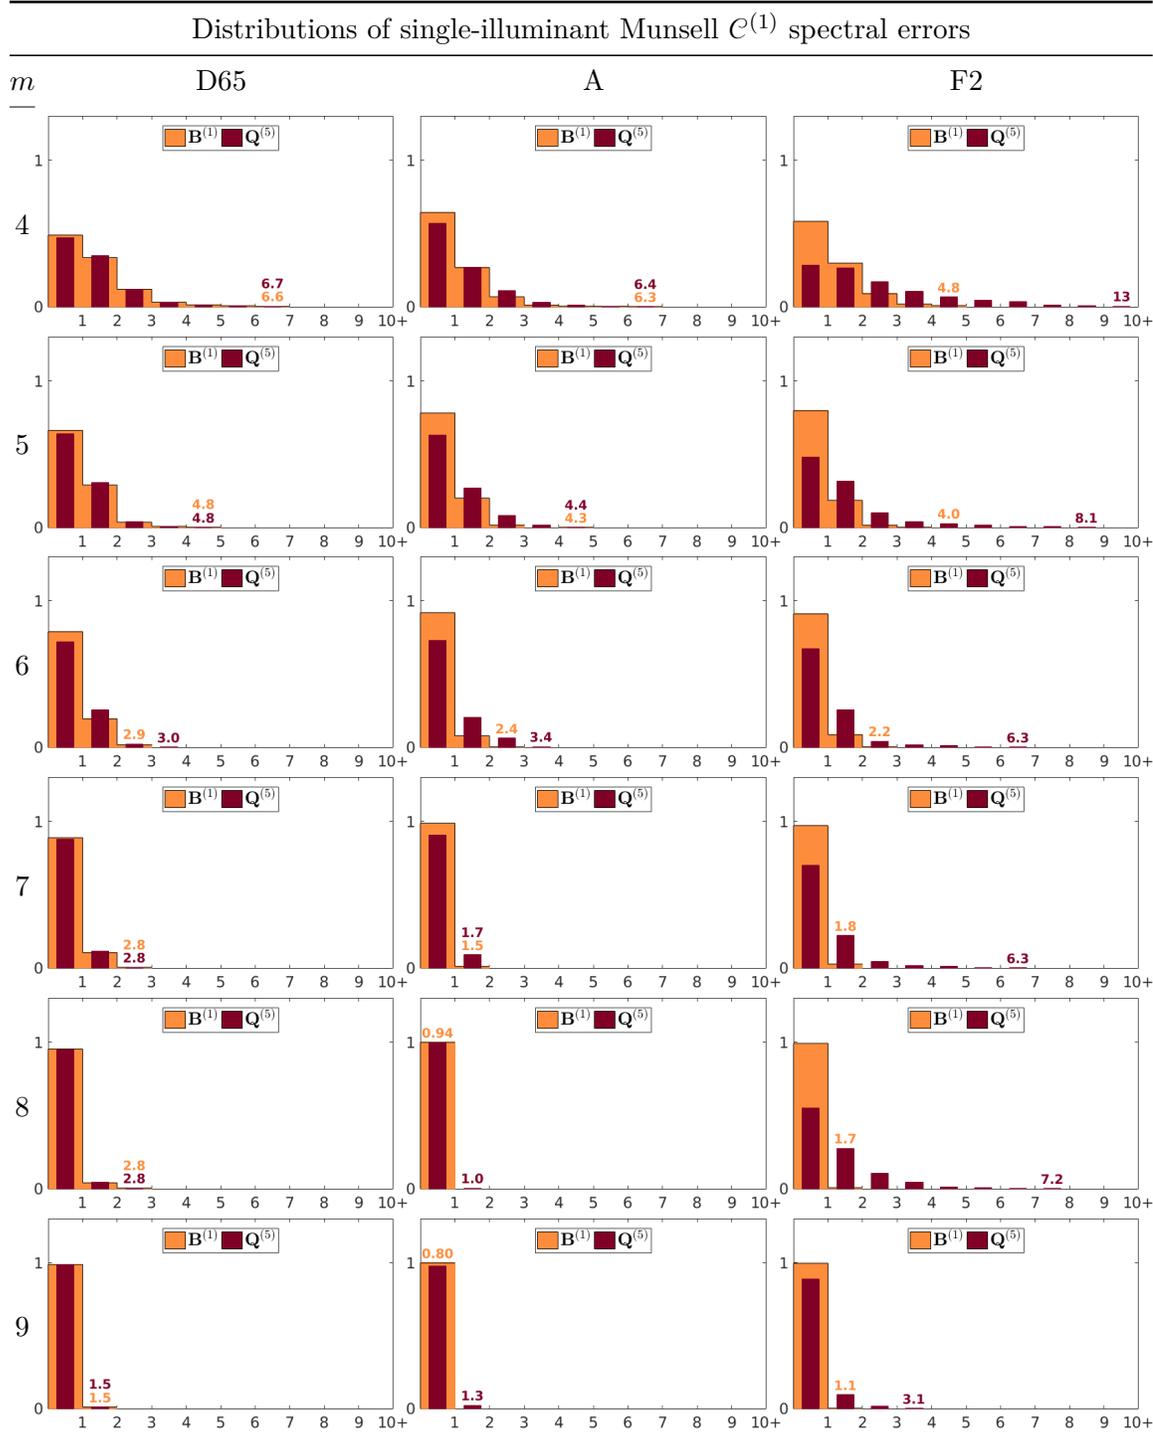


Figure A1: Distributions of NRMS spectral errors (as fractions of the totals) incurred when computing single-illuminant Munsell $C^{(1)}$ color signals by multiplying coefficients of $Q^{(5)}$, for dimensionalities $m \equiv 4, \dots, 9$. For comparison, errors incurred by projection onto $B^{(1)}$ are also shown. Maximum values, rounded to two significant figures, are shown in corresponding colors.

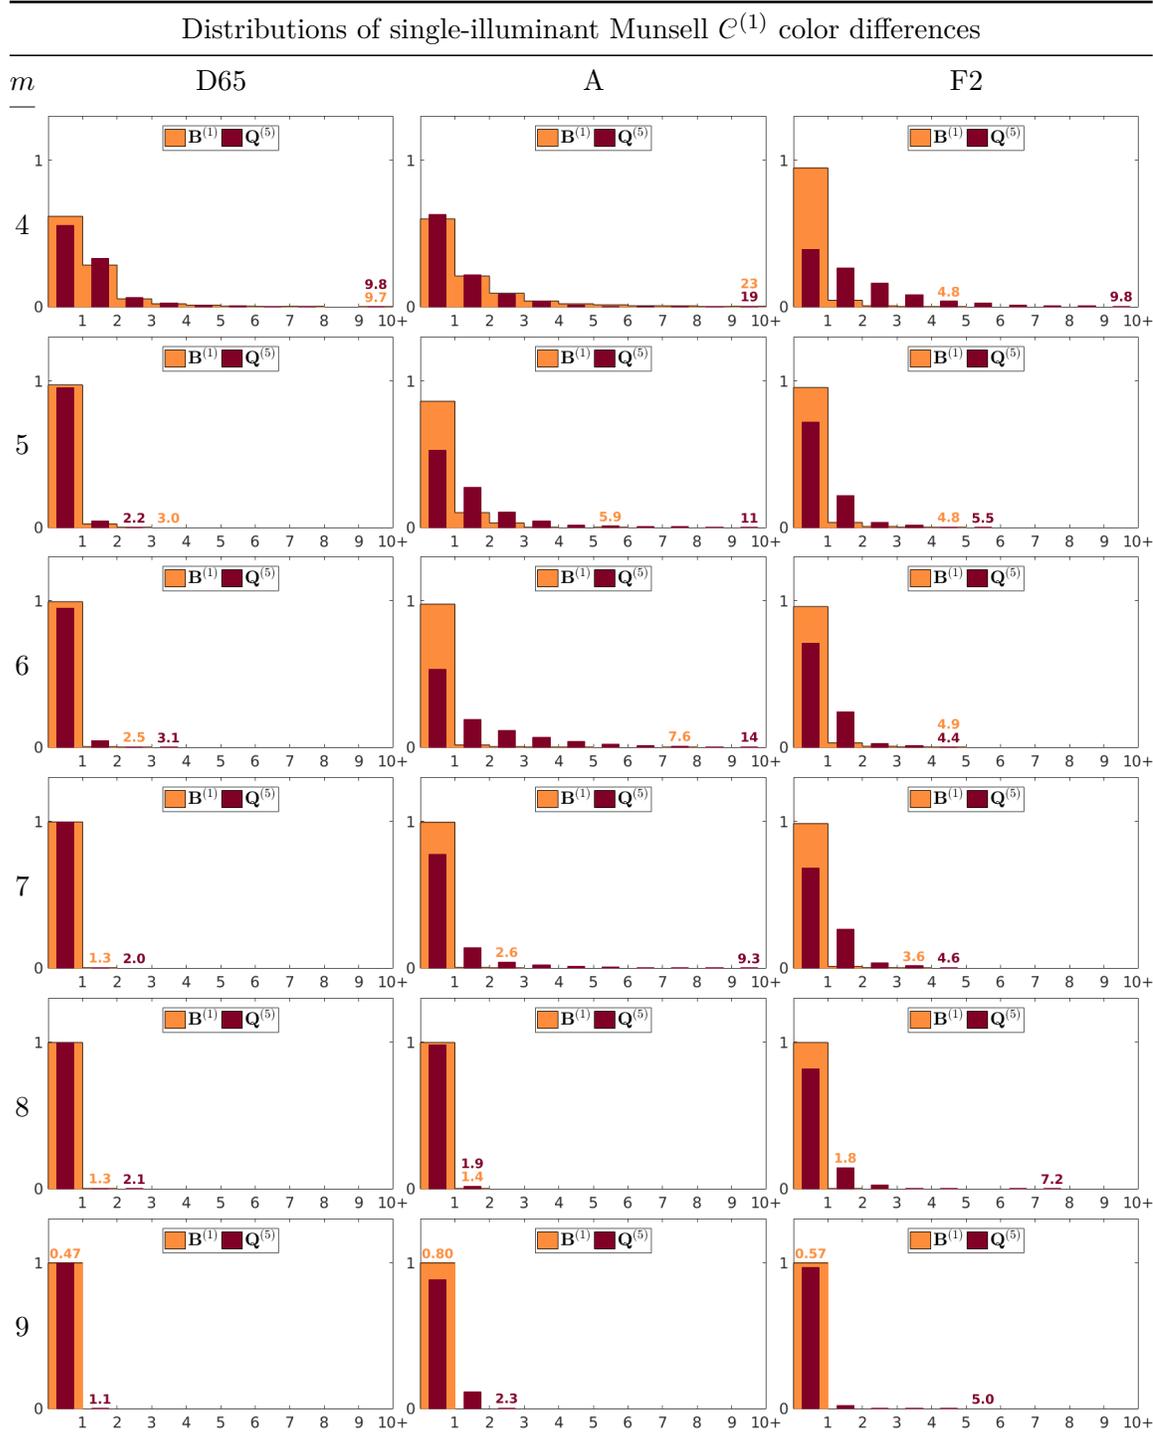


Figure A2: Distributions of ΔE_{00} color differences (as fractions of the totals) incurred when computing single-illuminant Munsell $\mathcal{C}^{(1)}$ color signals by multiplying coefficients of $\mathbf{Q}^{(5)}$, for dimensionalities $m \equiv 4, \dots, 9$. For comparison, errors incurred by projection onto $\mathbf{B}^{(1)}$ are also shown. Maximum values, rounded to two significant figures, are shown in corresponding colors.

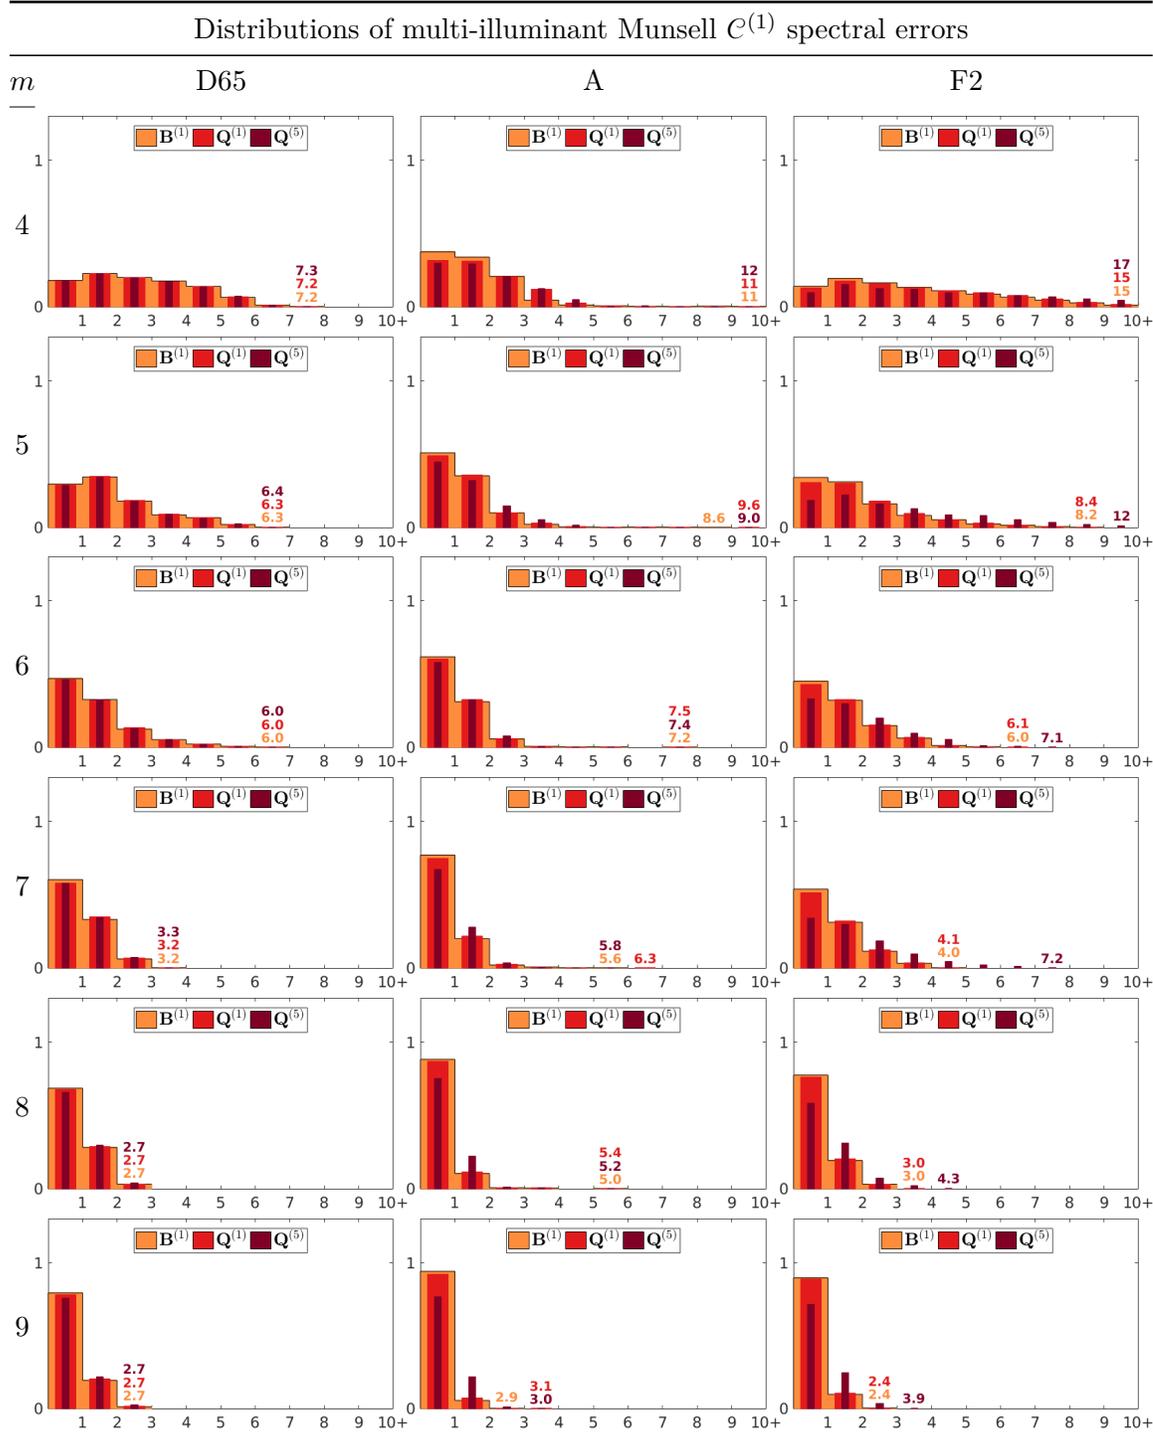


Figure A3: Distributions of NRMS spectral errors (as fractions of the totals) incurred when computing multi-illuminant Munsell $C^{(1)}$ color signals by multiplying coefficients of $Q^{(5)}$, for dimensionalities $m \equiv 4, \dots, 9$. For comparison, errors incurred by projection onto $B^{(1)}$ are also shown. Maximum values, rounded to two significant figures, are shown in corresponding colors.

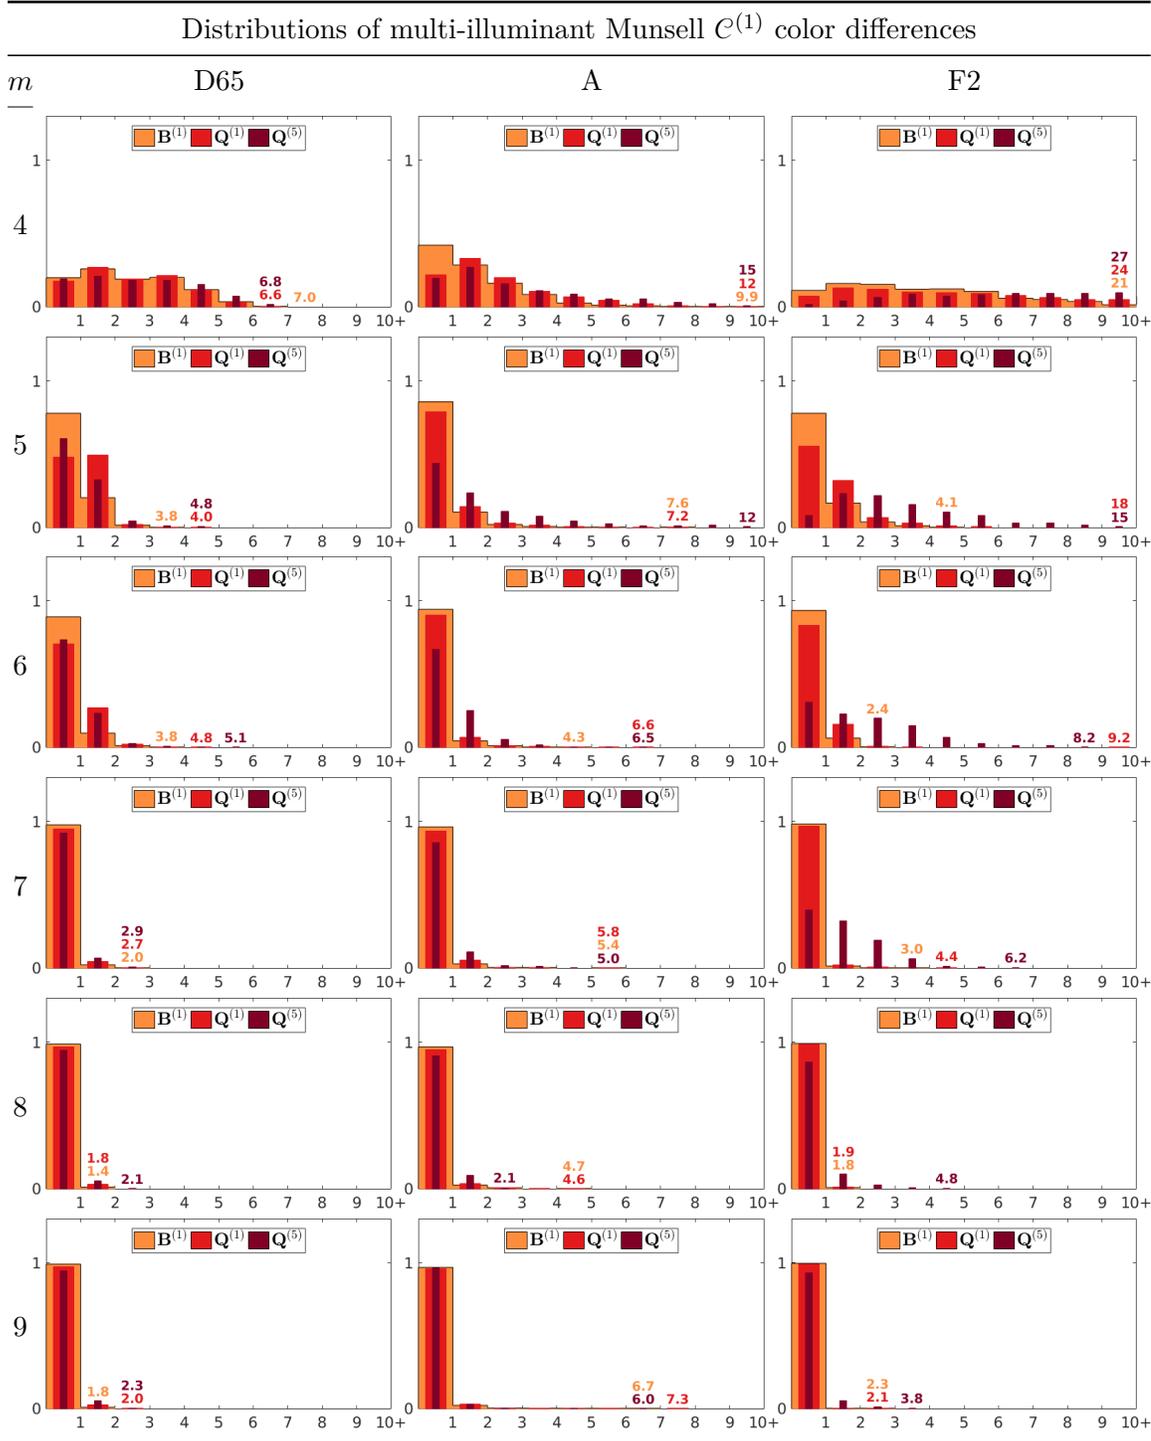


Figure A4: Distributions of ΔE_{00} color differences (as fractions of the totals) incurred when computing multi-illuminant Munsell $\mathcal{C}^{(1)}$ color signals by multiplying coefficients of $\mathbf{Q}^{(5)}$, for dimensionalities $m \equiv 4, \dots, 9$. For comparison, errors incurred by projection onto $\mathbf{B}^{(1)}$ are also shown. Maximum values, rounded to two significant figures, are shown in corresponding colors.

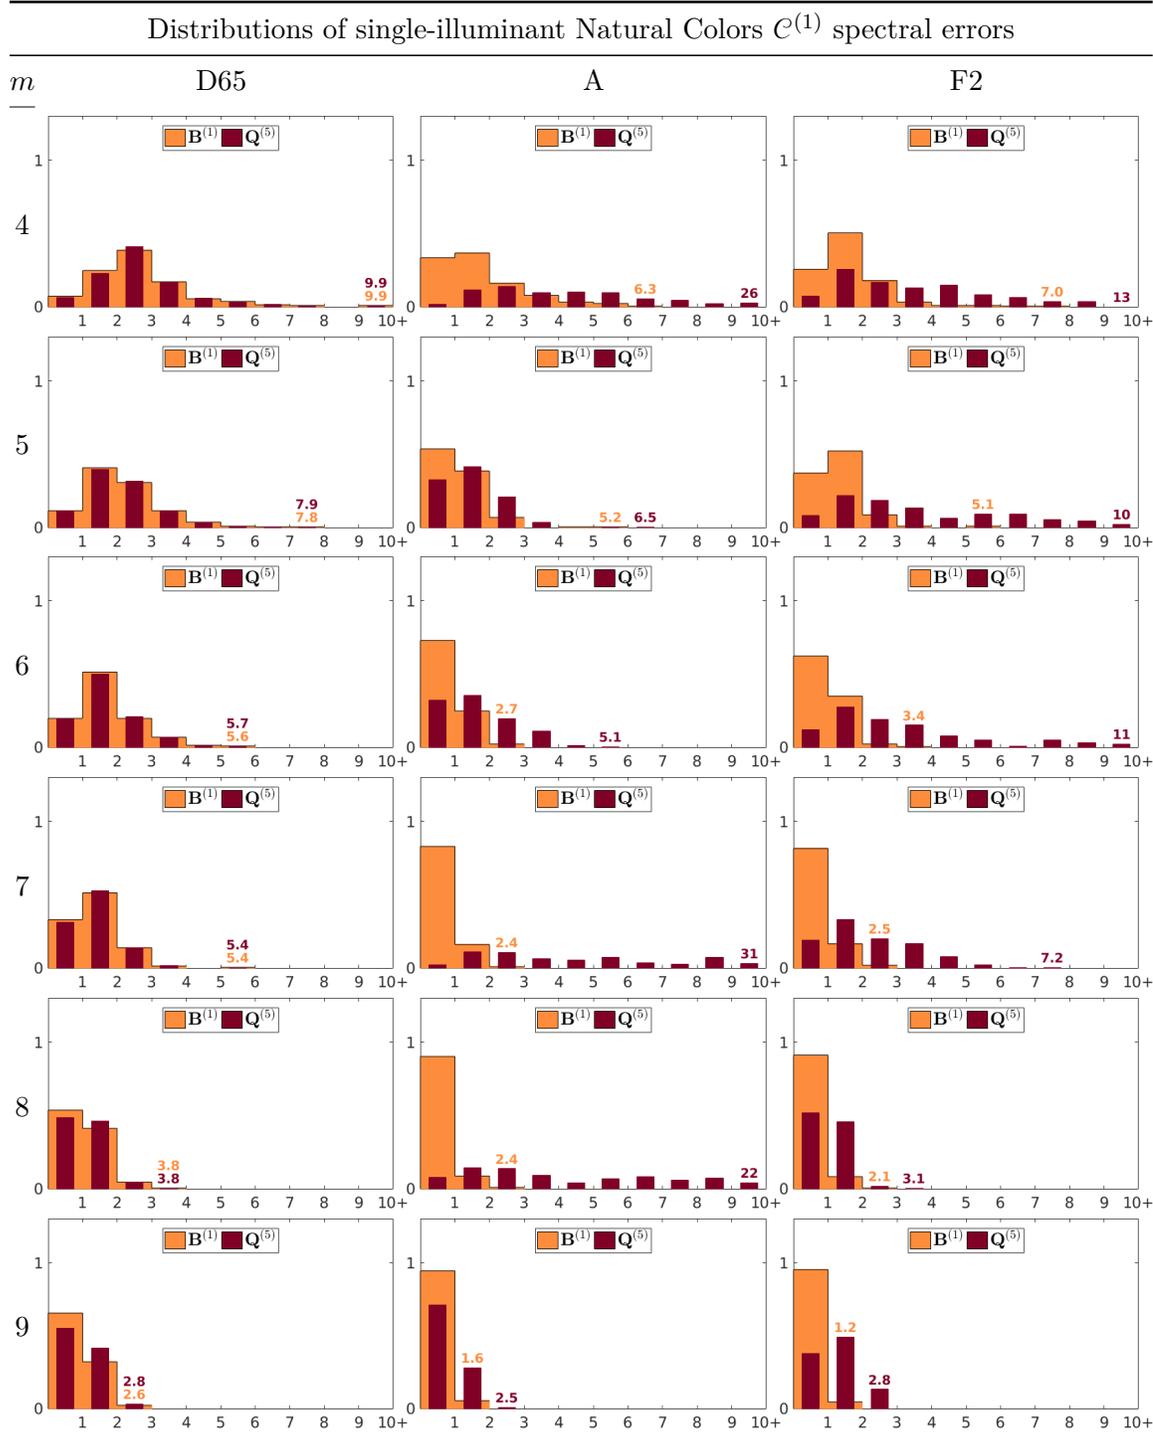


Figure A5: Distributions of NRMS spectral errors (as fractions of the totals) incurred when computing single-illuminant Natural Colors $\mathcal{C}^{(1)}$ color signals by multiplying coefficients of $\mathbf{Q}^{(5)}$, for dimensionalities $m \equiv 4, \dots, 9$. For comparison, errors incurred by projection onto $\mathbf{B}^{(1)}$ are also shown. Maximum values, rounded to two significant figures, are shown in corresponding colors.

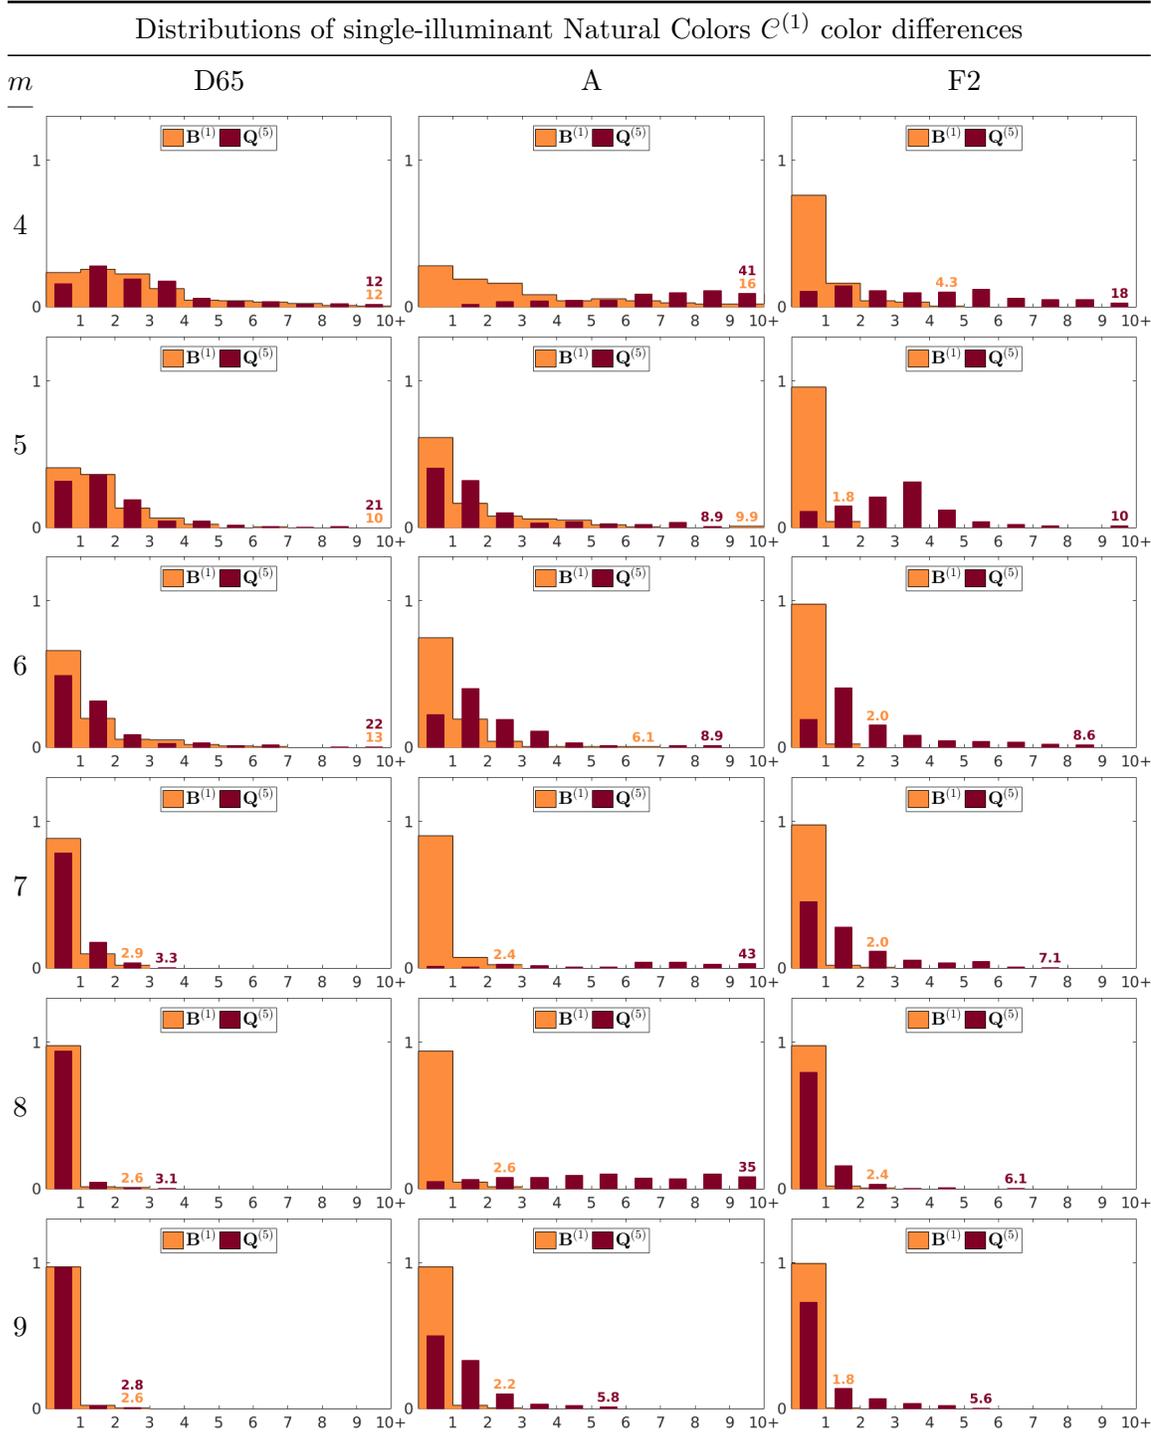


Figure A6: Distributions of ΔE_{00} color differences (as fractions of the totals) incurred when computing single-illuminant Natural Colors $\mathcal{C}^{(1)}$ color signals by multiplying coefficients of $Q^{(5)}$, for dimensionalities $m \equiv 4, \dots, 9$. For comparison, errors incurred by projection onto $B^{(1)}$ are also shown. Maximum values, rounded to two significant figures, are shown in corresponding colors.

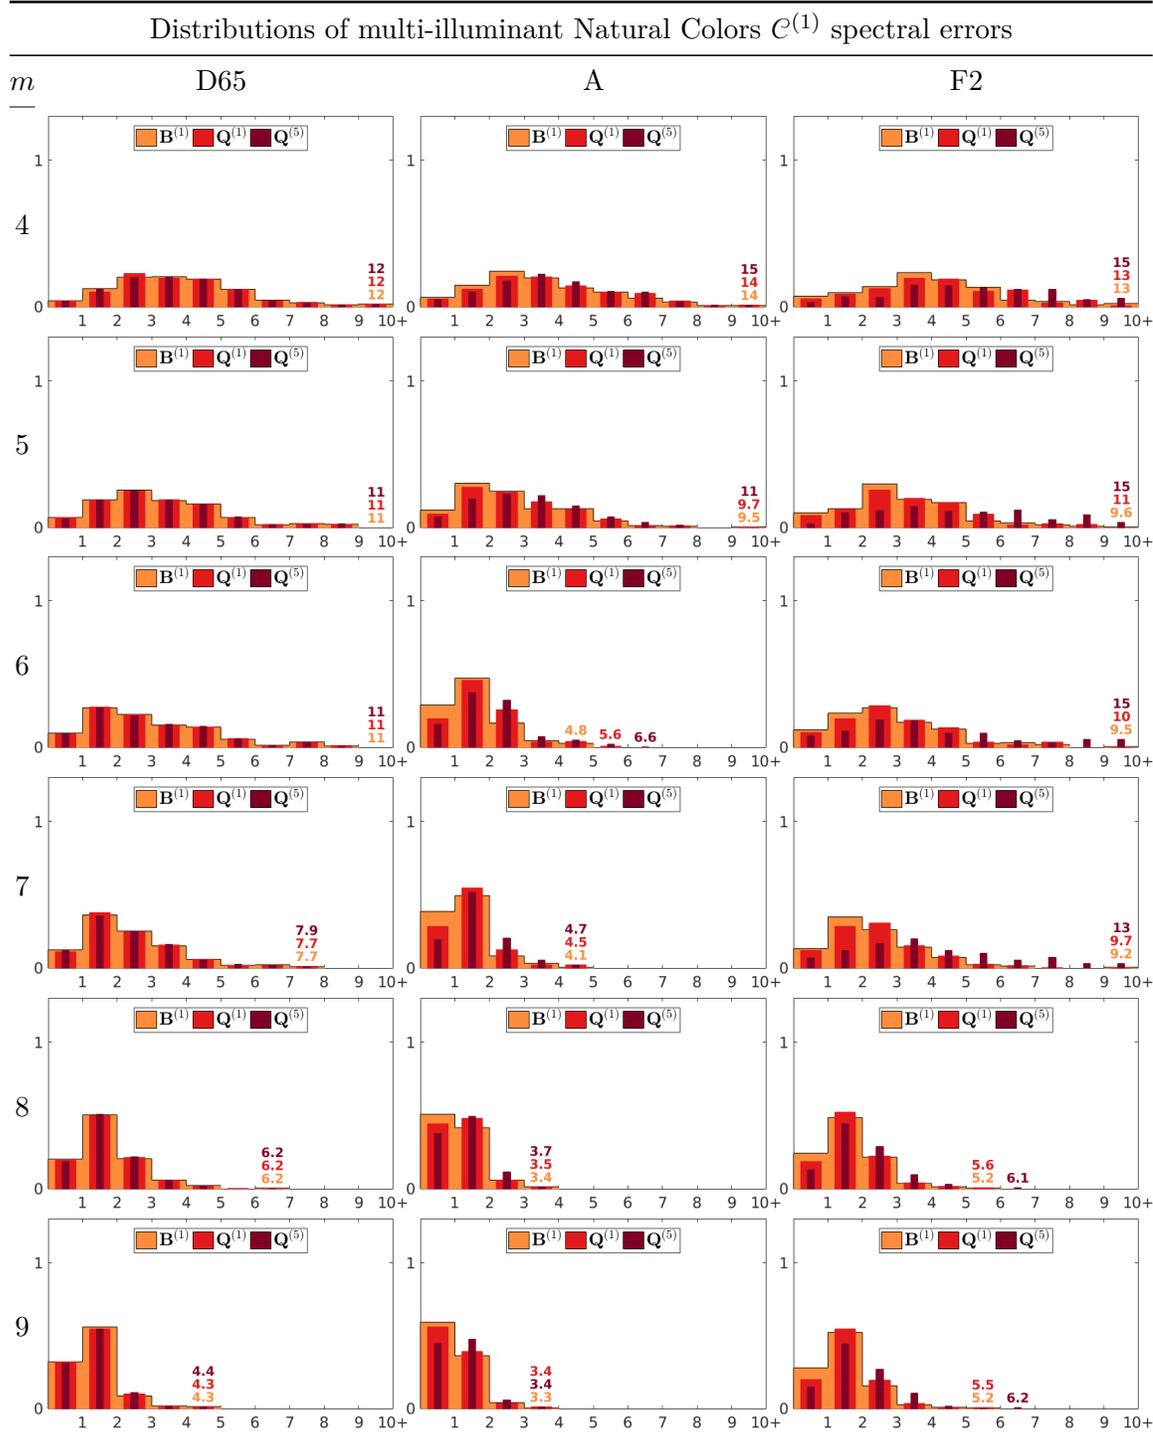


Figure A7: Distributions of NRMS spectral errors (as fractions of the totals) incurred when computing multi-illuminant Natural Colors $\mathcal{C}^{(1)}$ color signals by multiplying coefficients of $Q^{(5)}$, for dimensionalities $m \equiv 4, \dots, 9$. For comparison, errors incurred by projection onto $B^{(1)}$ are also shown. Maximum values, rounded to two significant figures, are shown in corresponding colors.

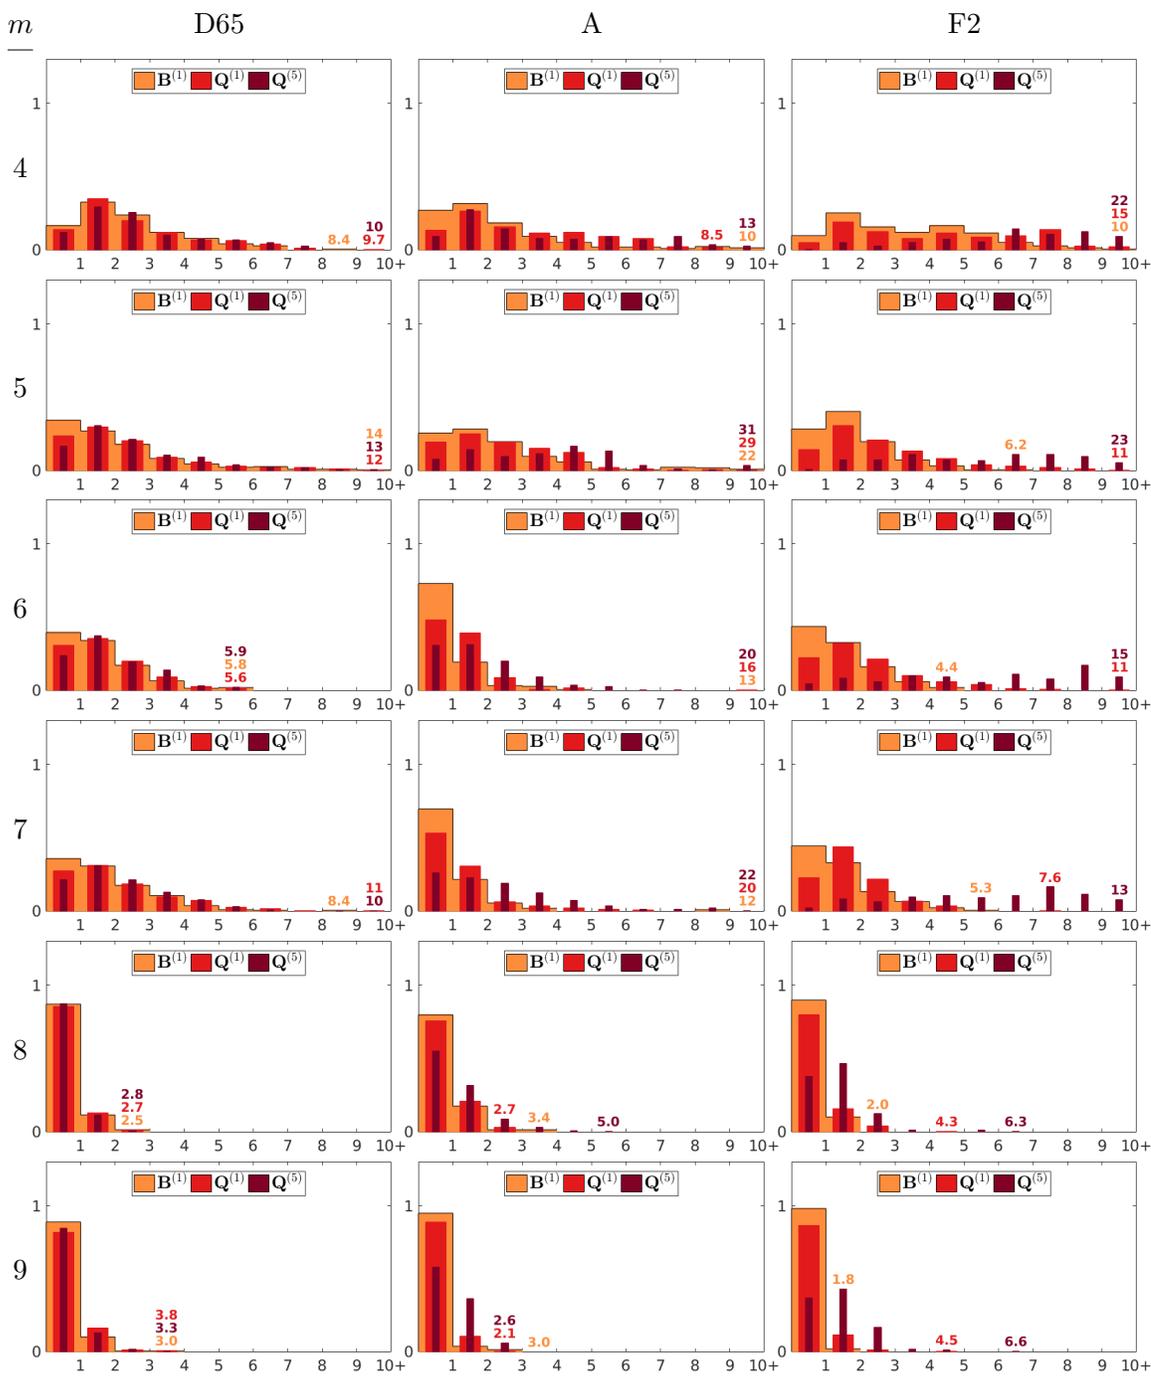
Distributions of multi-illuminant Natural Colors $\mathcal{C}^{(1)}$ color differences

Figure A8: Distributions of ΔE_{00} color differences (as fractions of the totals) incurred when computing multi-illuminant Natural Colors $\mathcal{C}^{(1)}$ color signals by multiplying coefficients of $\mathbf{Q}^{(5)}$, for dimensionalities $m \equiv 4, \dots, 9$. For comparison, errors incurred by projection onto $\mathbf{B}^{(1)}$ are also shown. Maximum values, rounded to two significant figures, are shown in corresponding colors.

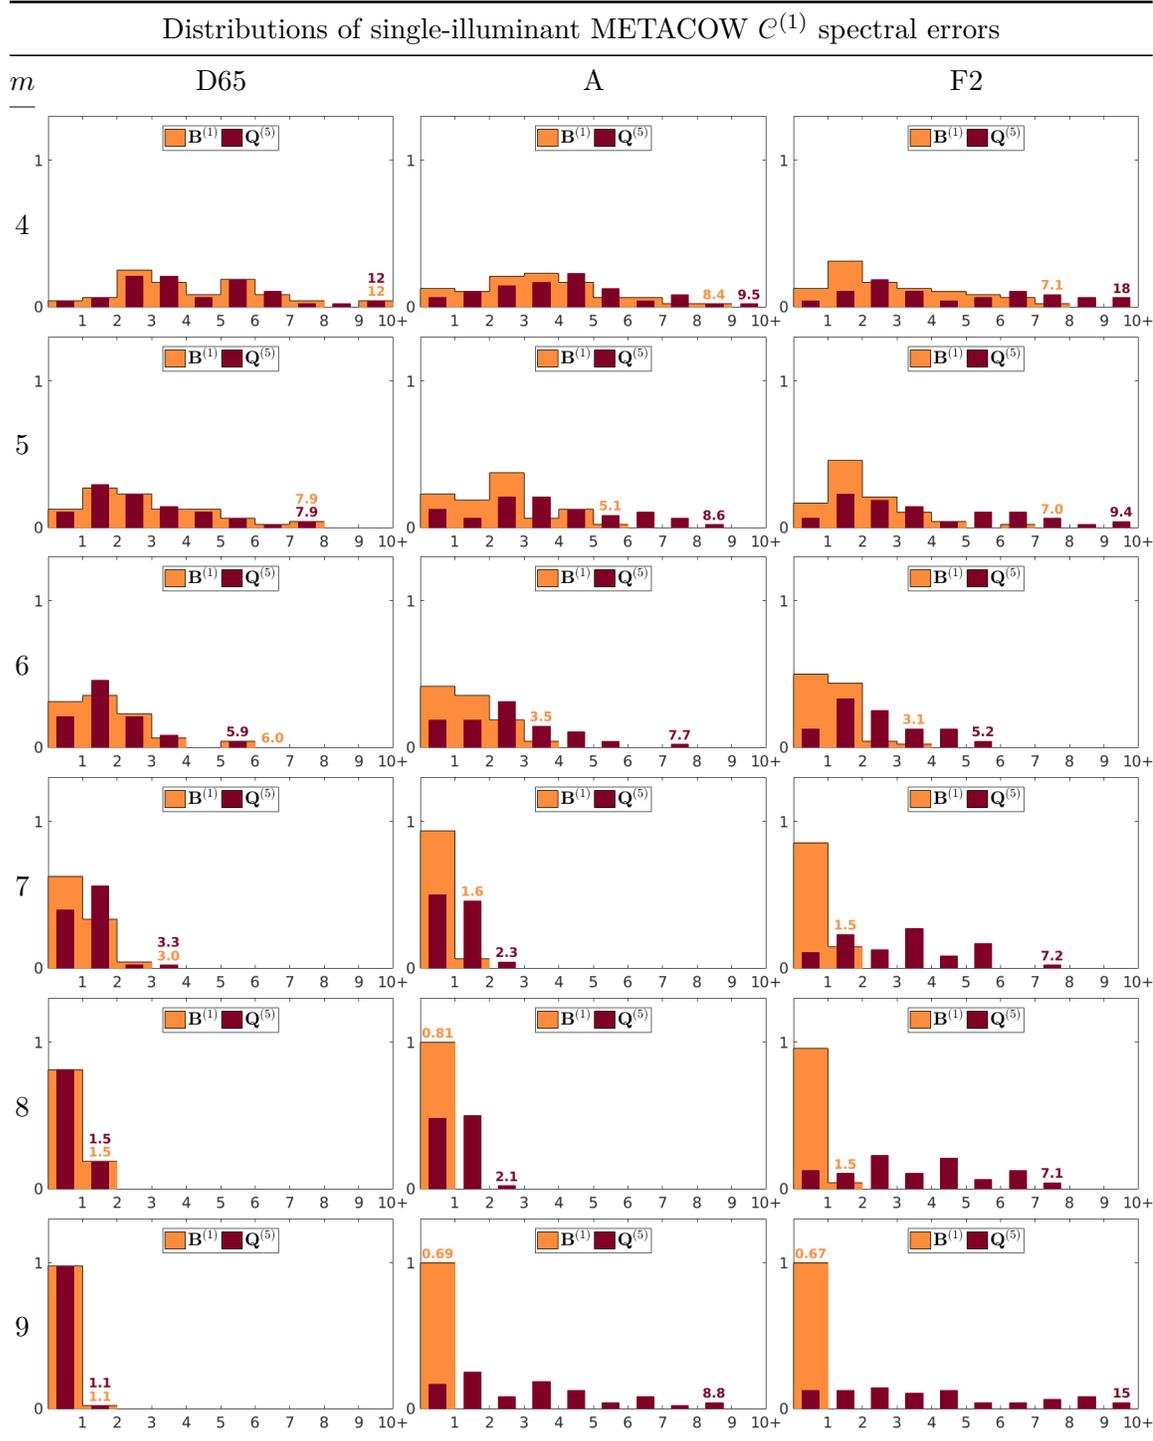


Figure A9: Distributions of NRMS spectral errors (as fractions of the totals) incurred when computing single-illuminant METACOW $c^{(1)}$ color signals by multiplying coefficients of $\mathbf{Q}^{(5)}$, for dimensionalities $m \equiv 4, \dots, 9$. For comparison, errors incurred by projection onto $\mathbf{B}^{(1)}$ are also shown. Maximum values, rounded to two significant figures, are shown in corresponding colors.

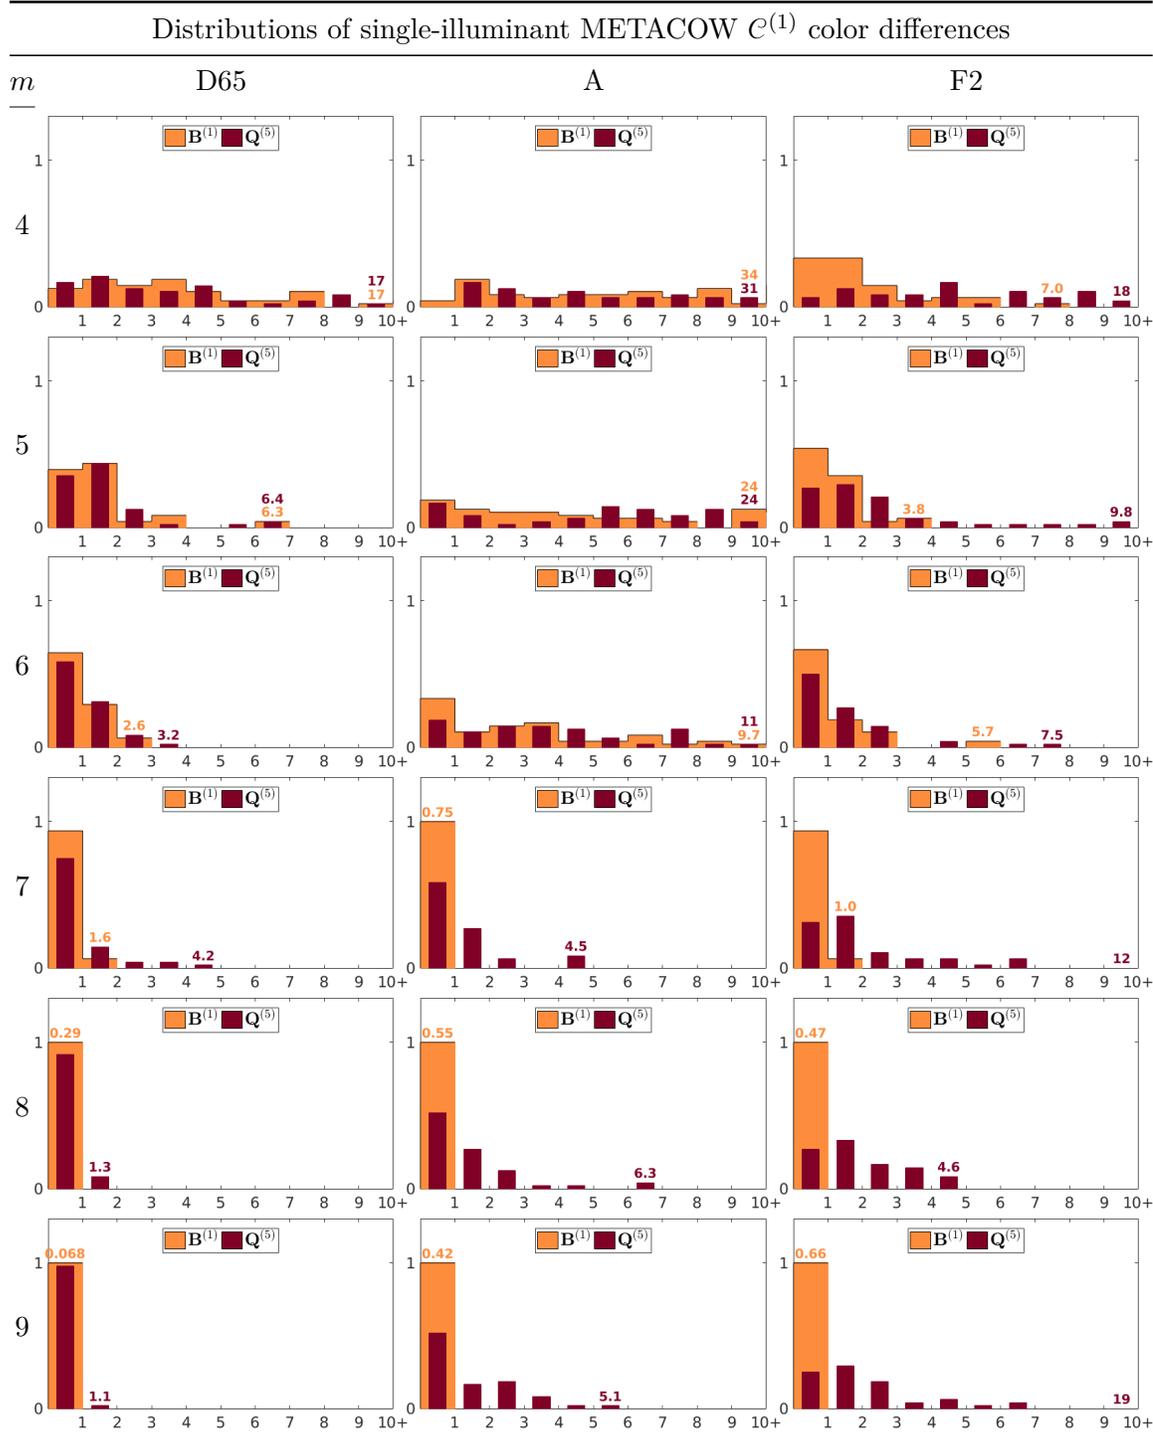


Figure A10: Distributions of ΔE_{00} color differences (as fractions of the totals) incurred when computing single-illuminant METACOW $\mathcal{C}^{(1)}$ color signals by multiplying coefficients of $\mathbf{Q}^{(5)}$, for dimensionalities $m \equiv 4, \dots, 9$. For comparison, errors incurred by projection onto $\mathbf{B}^{(1)}$ are also shown. Maximum values, rounded to two significant figures, are shown in corresponding colors.

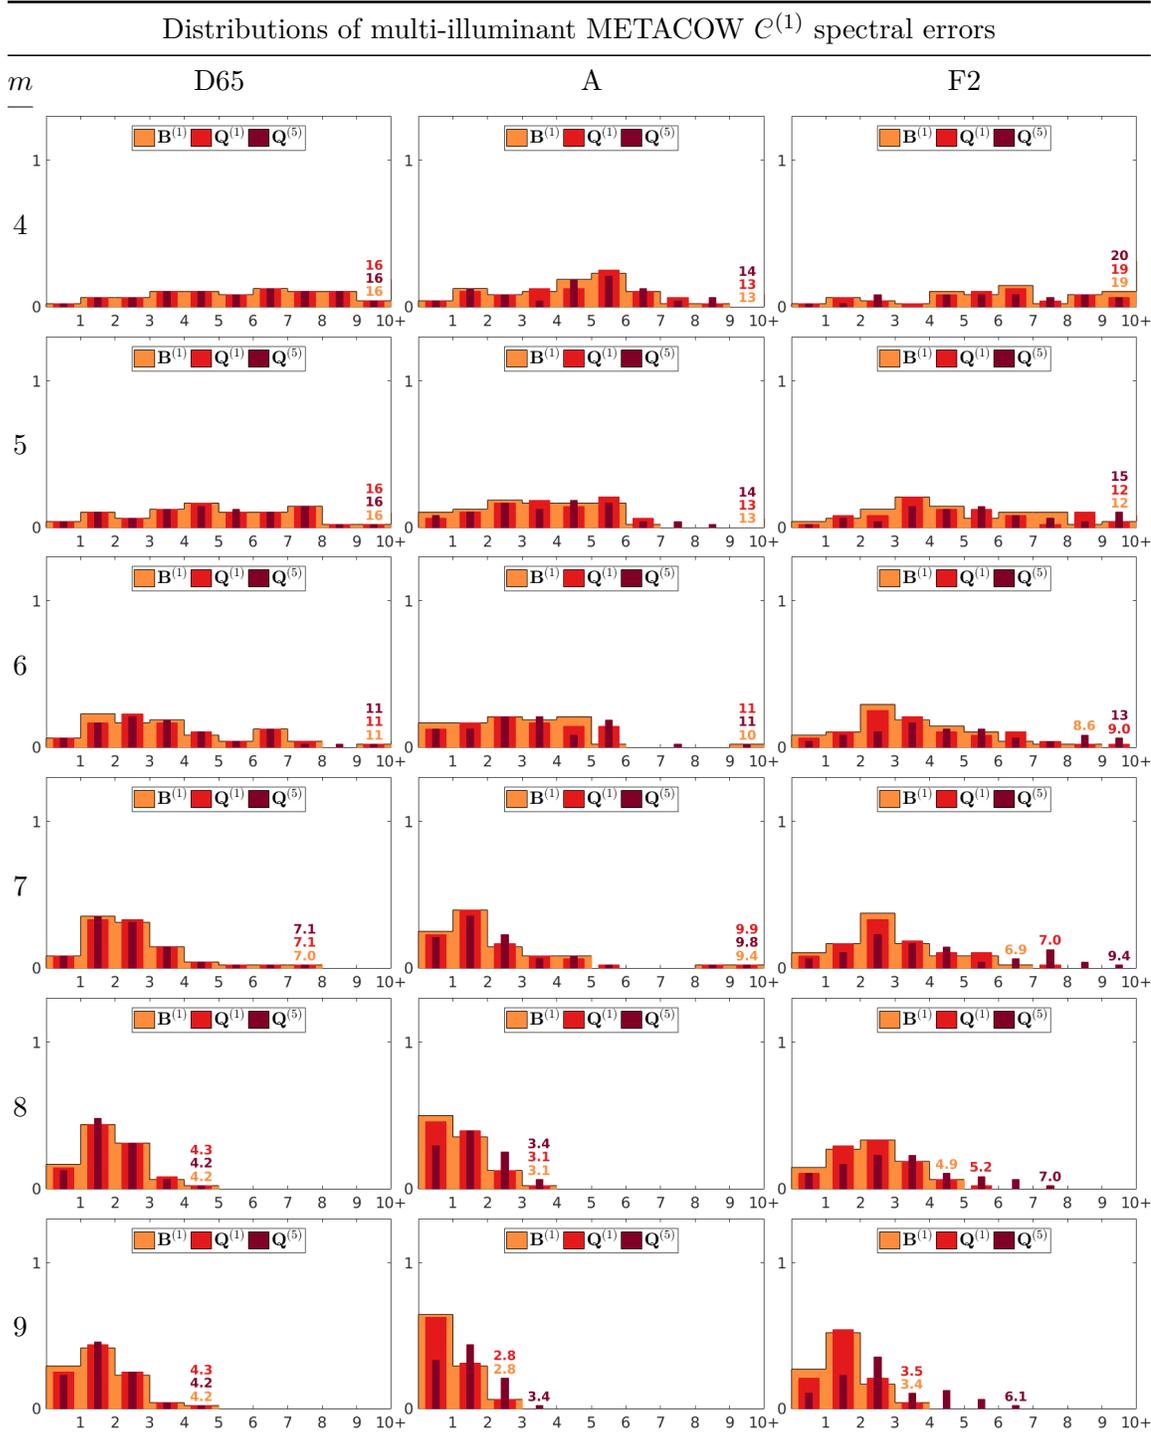


Figure A11: Distributions of NRMS spectral errors (as fractions of the totals) incurred when computing multi-illuminant METACOW $c^{(1)}$ color signals by multiplying coefficients of $Q^{(5)}$, for dimensionalities $m \equiv 4, \dots, 9$. For comparison, errors incurred by projection onto $B^{(1)}$ are also shown. Maximum values, rounded to two significant figures, are shown in corresponding colors.

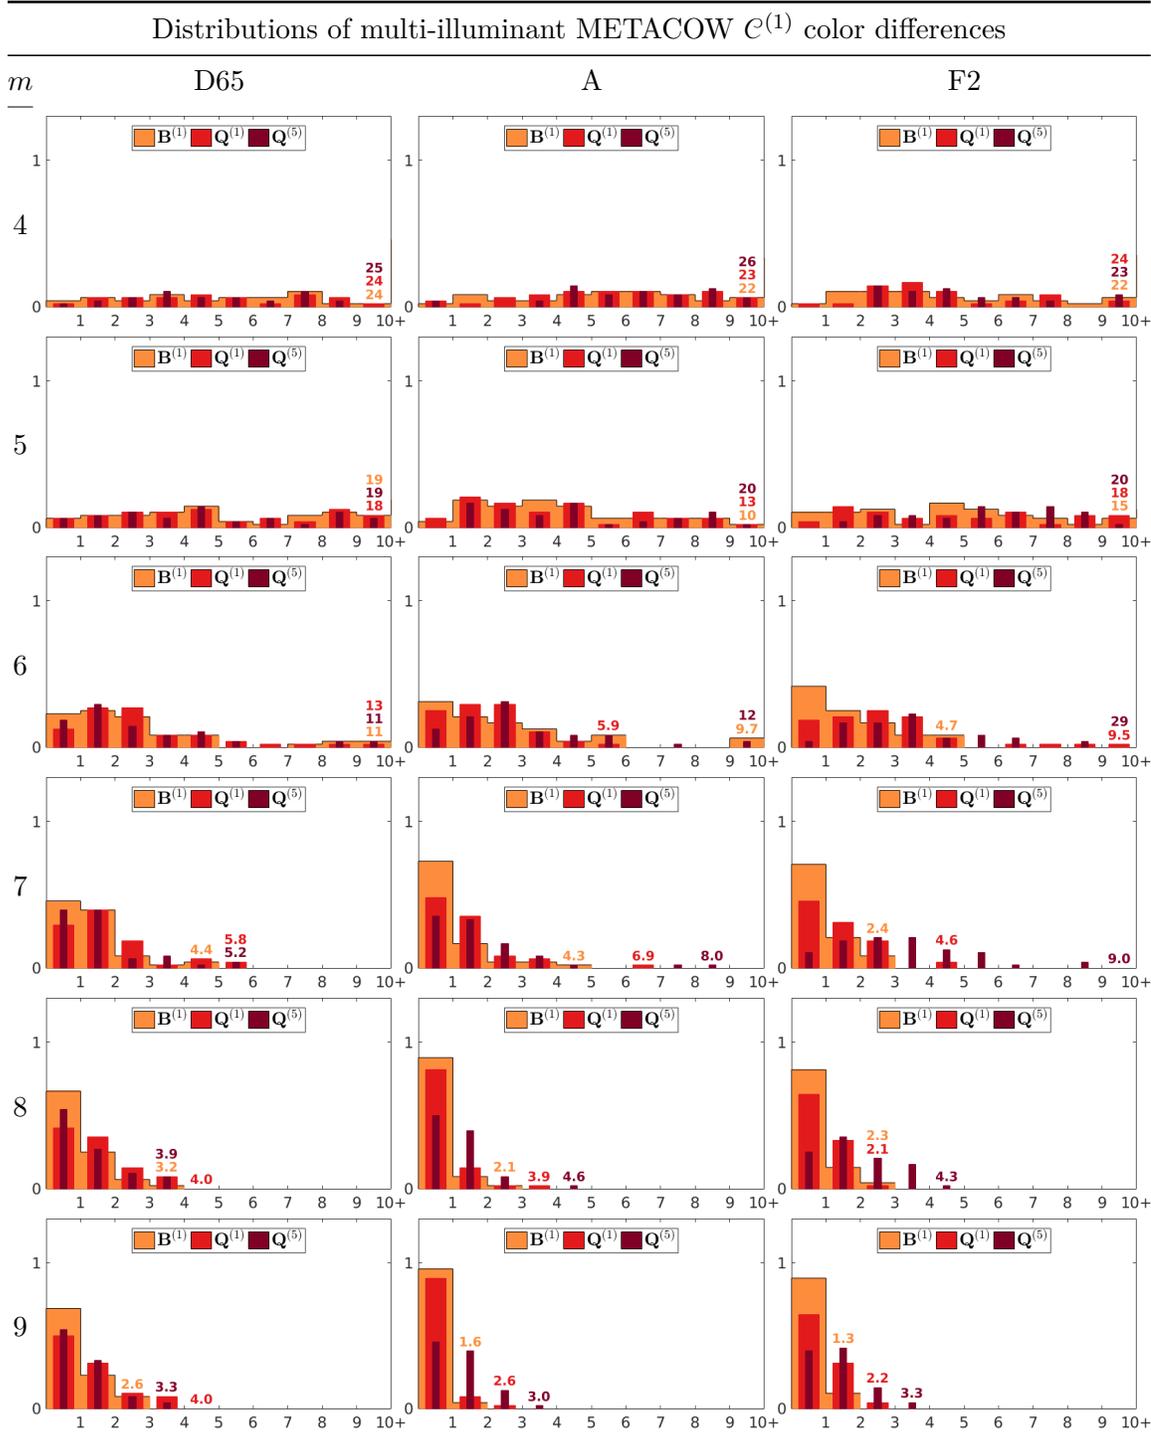


Figure A12: Distributions of ΔE_{00} color differences (as fractions of the totals) incurred when computing multi-illuminant METACOW $\mathcal{C}^{(1)}$ color signals by multiplying coefficients of $\mathbf{Q}^{(5)}$, for dimensionalities $m \equiv 4, \dots, 9$. For comparison, errors incurred by projection onto $\mathbf{B}^{(1)}$ are also shown. Maximum values, rounded to two significant figures, are shown in corresponding colors.

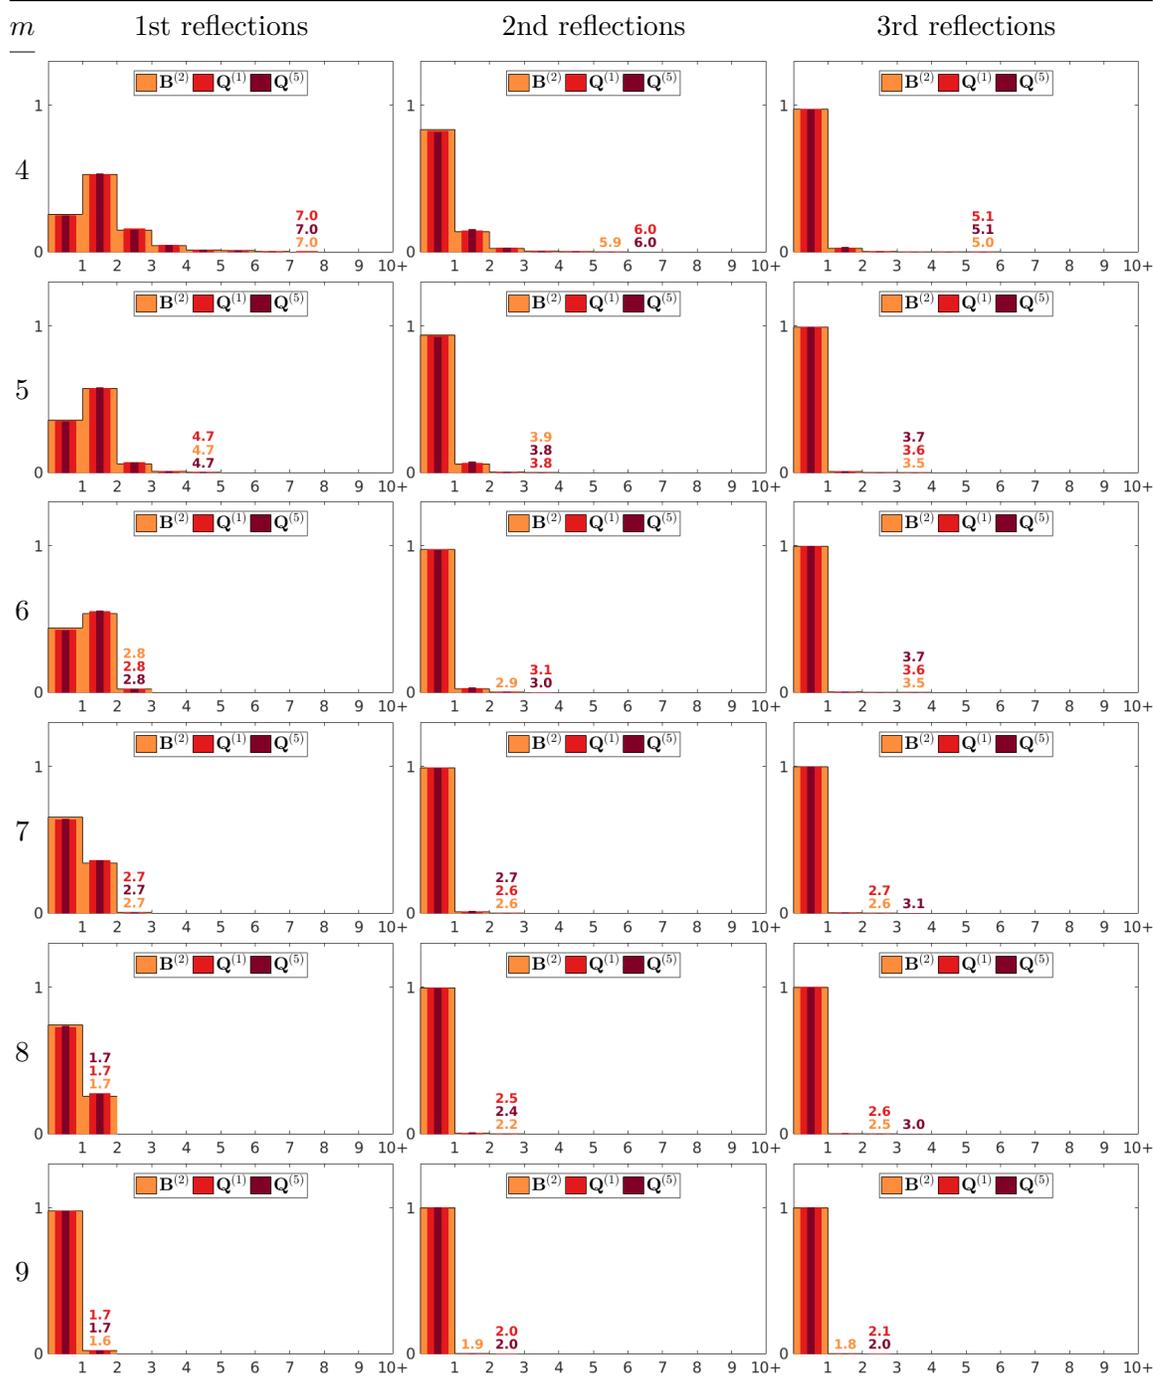
Distributions of Munsell/D65 $C^{(3)}$ spectral errors

Figure A13: Distributions of NRMS spectral errors (as fractions of the totals) incurred when computing Munsell/D65 $C^{(3)}$ color signals by multiplying coefficients of $Q^{(1)}$ or $Q^{(5)}$, for dimensionalities $m \equiv 4, \dots, 9$ and first, second and third reflections. For comparison, errors incurred by projection onto $B^{(2)}$ are also shown. Maximum values, rounded to two significant figures, are shown in corresponding colors.

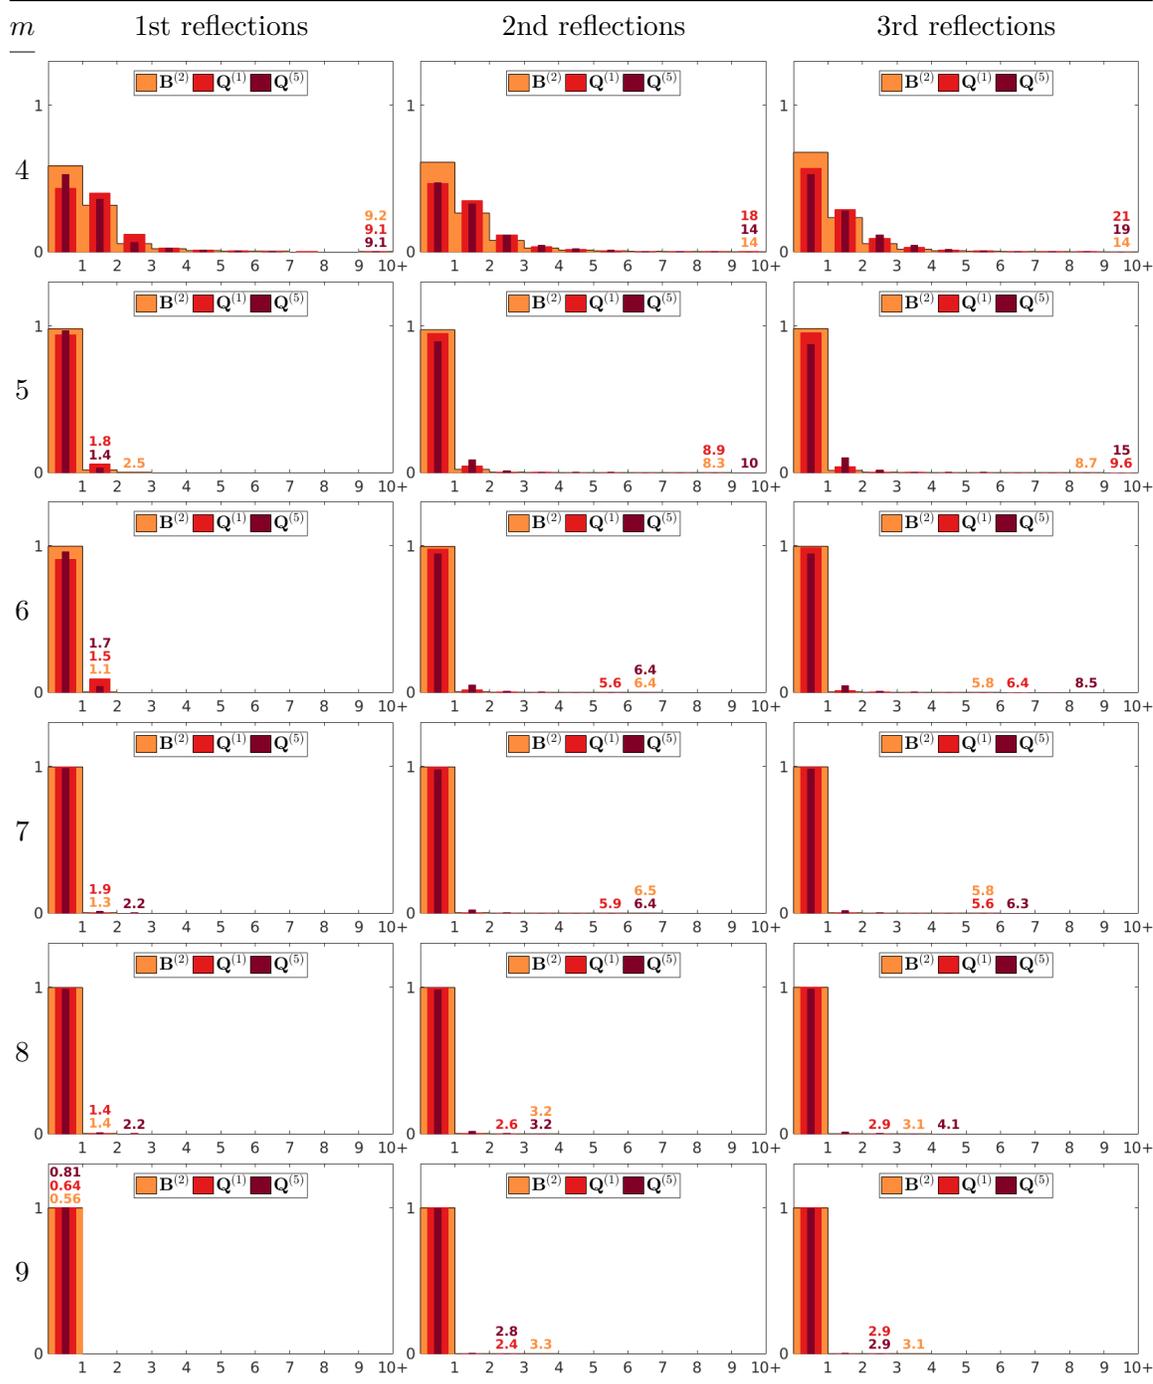
Distributions of Munsell/D65 $C^{(3)}$ color differences

Figure A14: Distributions of ΔE_{00} color differences (as fractions of the totals) incurred when computing Munsell/D65 $C^{(3)}$ color signals by multiplying coefficients of $Q^{(1)}$ or $Q^{(5)}$, for dimensionalities $m \equiv 4, \dots, 9$ and first, second and third reflections. For comparison, differences incurred by projection onto $B^{(2)}$ are also shown. Maximum values, rounded to two significant figures, are shown in corresponding colors.

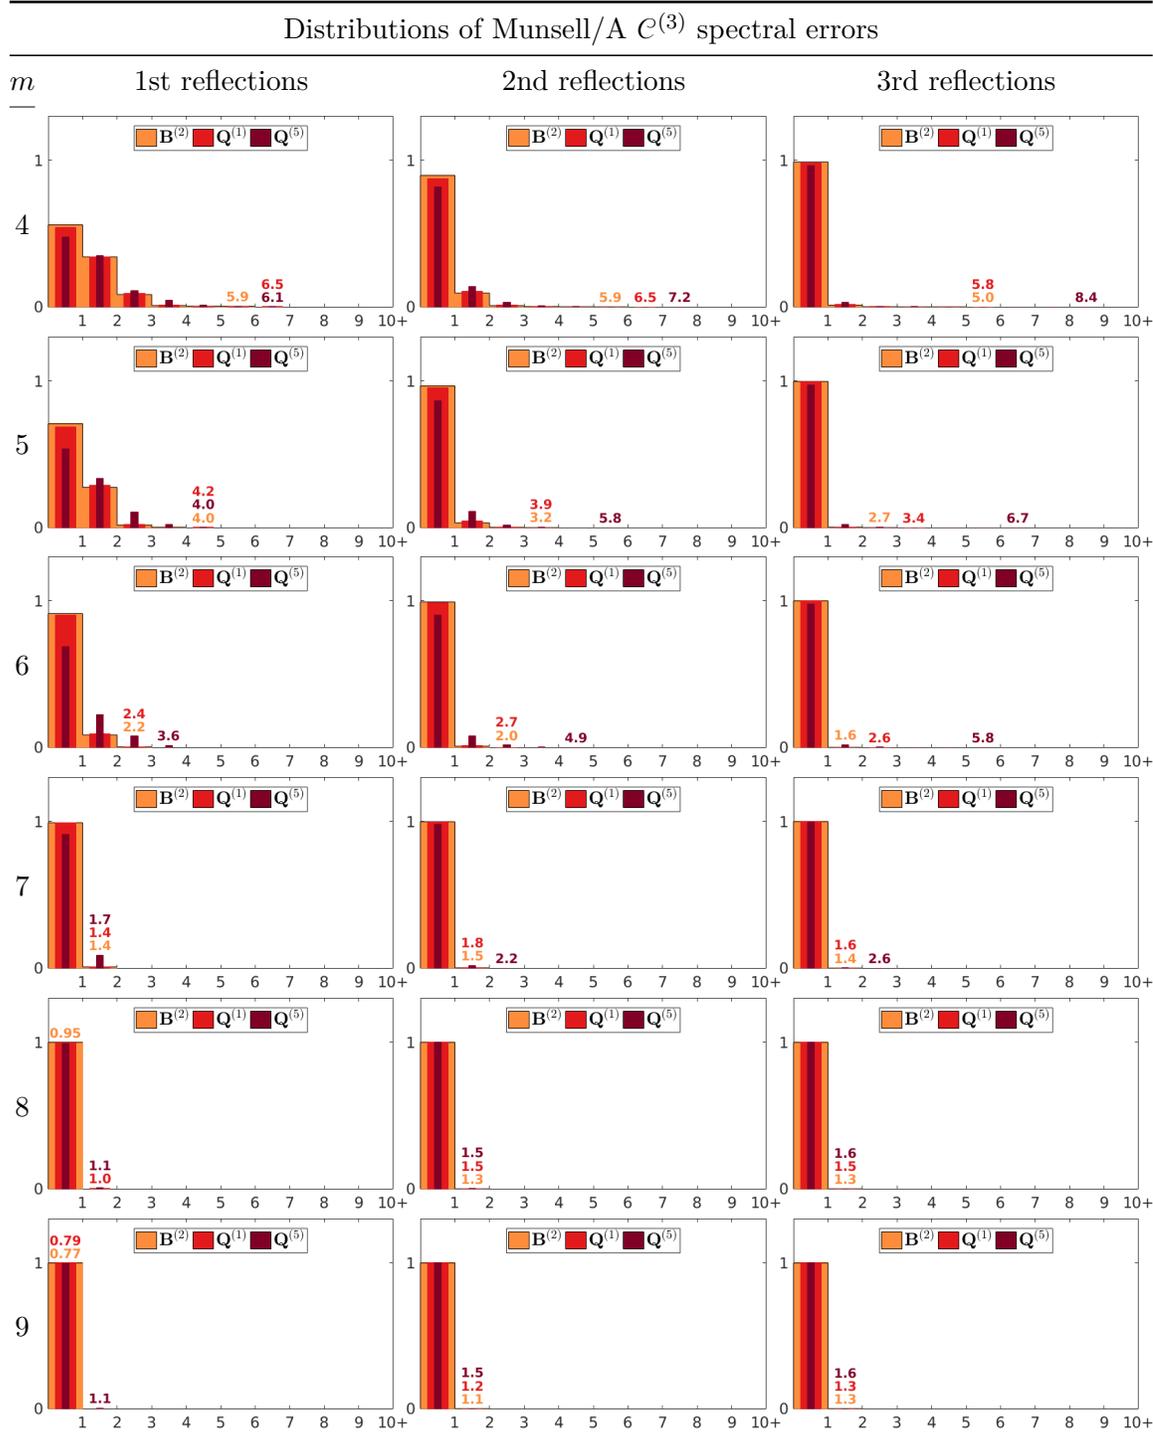


Figure A15: Distributions of NRMS spectral errors (as fractions of the totals) incurred when computing Munsell/A $C^{(3)}$ color signals by multiplying coefficients of $Q^{(1)}$ or $Q^{(5)}$, for dimensionalities $m \equiv 4, \dots, 9$ and first, second and third reflections. For comparison, errors incurred by projection onto $B^{(2)}$ are also shown. Maximum values, rounded to two significant figures, are shown in corresponding colors.

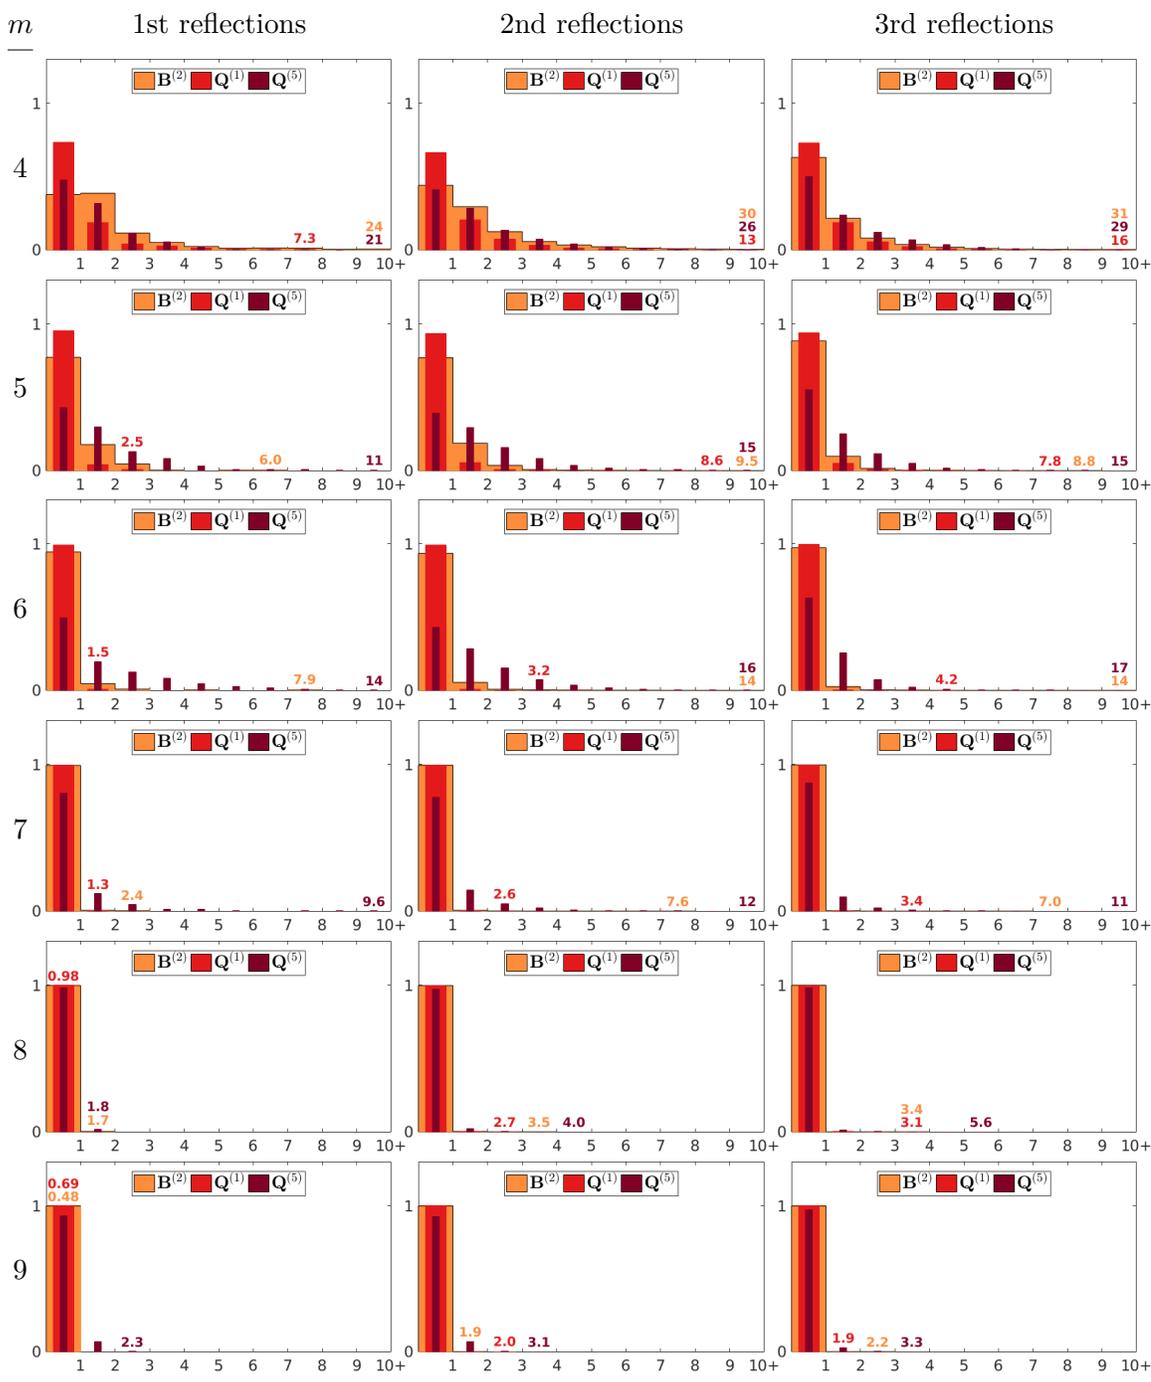
Distributions of Munsell/A $C^{(3)}$ color differences

Figure A16: Distributions of ΔE_{00} color differences (as fractions of the totals) incurred when computing Munsell/A $C^{(3)}$ color signals by multiplying coefficients of $Q^{(1)}$ or $Q^{(5)}$, for dimensionalities $m \equiv 4, \dots, 9$ and first, second and third reflections. For comparison, differences incurred by projection onto $B^{(2)}$ are also shown. Maximum values, rounded to two significant figures, are shown in corresponding colors.

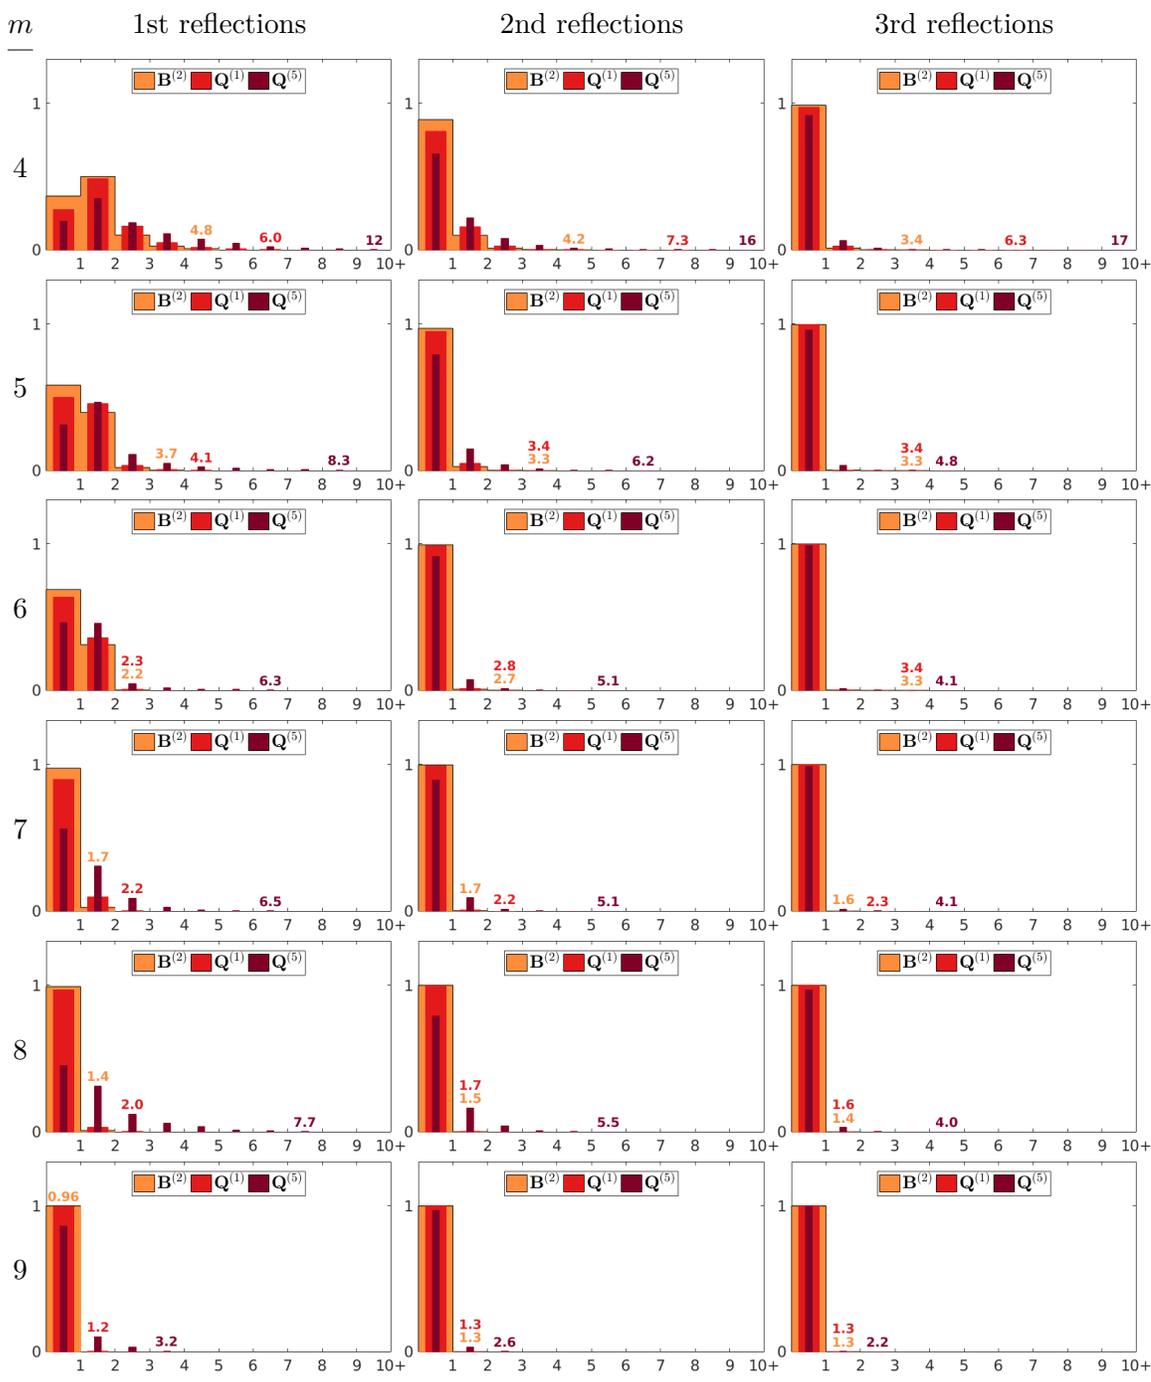
Distributions of Munsell/F2 $\mathcal{C}^{(3)}$ spectral errors

Figure A17: Distributions of NRMS spectral errors (as fractions of the totals) incurred when computing Munsell/F2 $\mathcal{C}^{(3)}$ color signals by multiplying coefficients of $\mathbf{Q}^{(1)}$ or $\mathbf{Q}^{(5)}$, for dimensionalities $m \equiv 4, \dots, 9$ and first, second and third reflections. For comparison, errors incurred by projection onto $\mathbf{B}^{(2)}$ are also shown. Maximum values, rounded to two significant figures, are shown in corresponding colors.

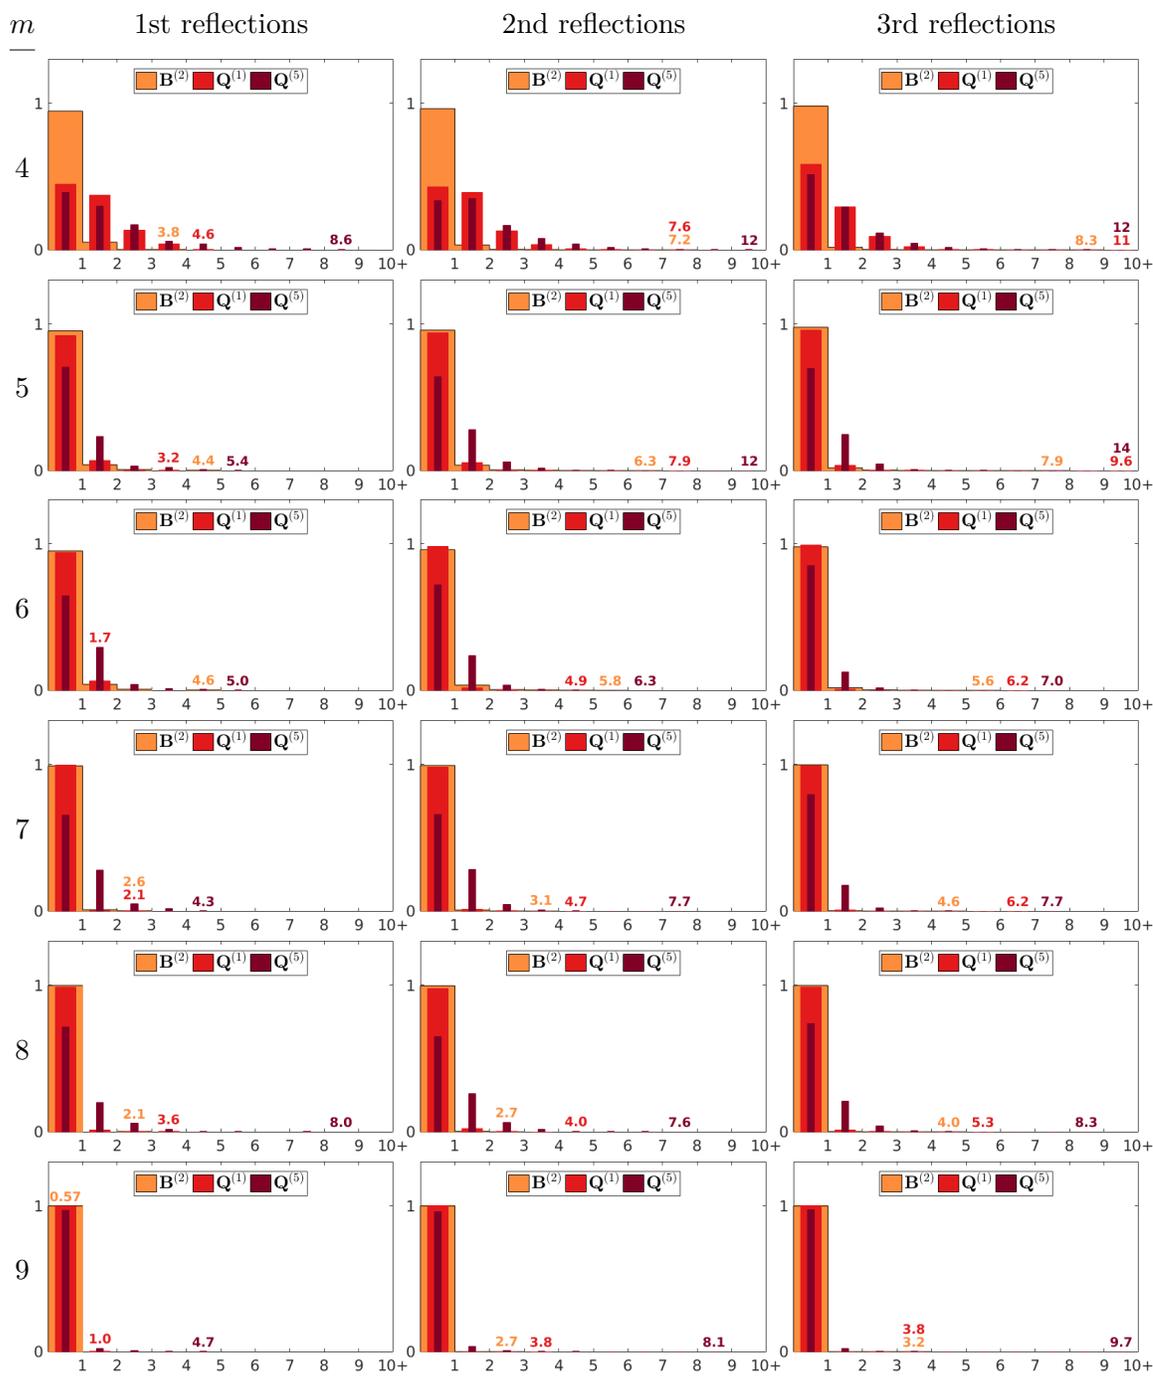
Distributions of Munsell/F2 $C^{(3)}$ color differences

Figure A18: Distributions of ΔE_{00} color differences (as fractions of the totals) incurred when computing Munsell/F2 $C^{(3)}$ color signals by multiplying coefficients of $Q^{(1)}$ or $Q^{(5)}$, for dimensionalities $m \equiv 4, \dots, 9$ and first, second and third reflections. For comparison, differences incurred by projection onto $B^{(2)}$ are also shown. Maximum values, rounded to two significant figures, are shown in corresponding colors.

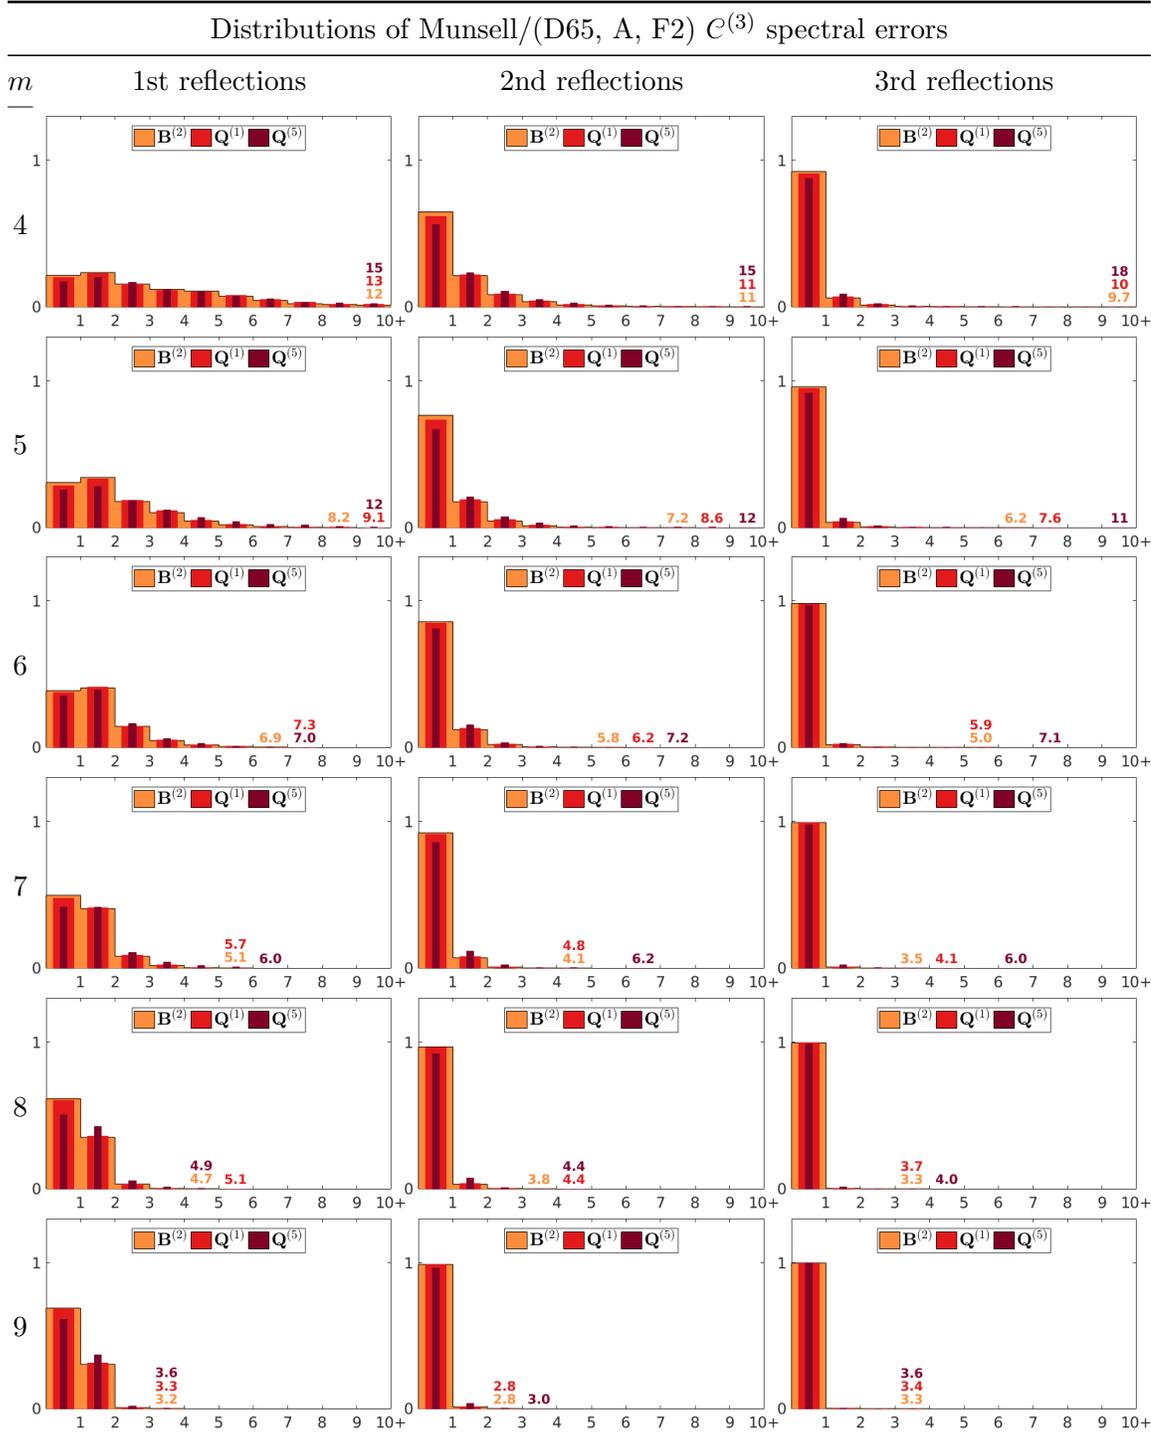


Figure A19: Distributions of NRMS spectral errors (as fractions of the totals) incurred when computing Munsell/(D65, A, F2) $C^{(3)}$ color signals by multiplying coefficients of $Q^{(1)}$ or $Q^{(5)}$, for dimensionalities $m \equiv 4, \dots, 9$ and first, second and third reflections. For comparison, errors incurred by projection onto $B^{(2)}$ are also shown. Maximum values, rounded to two significant figures, are shown in corresponding colors.

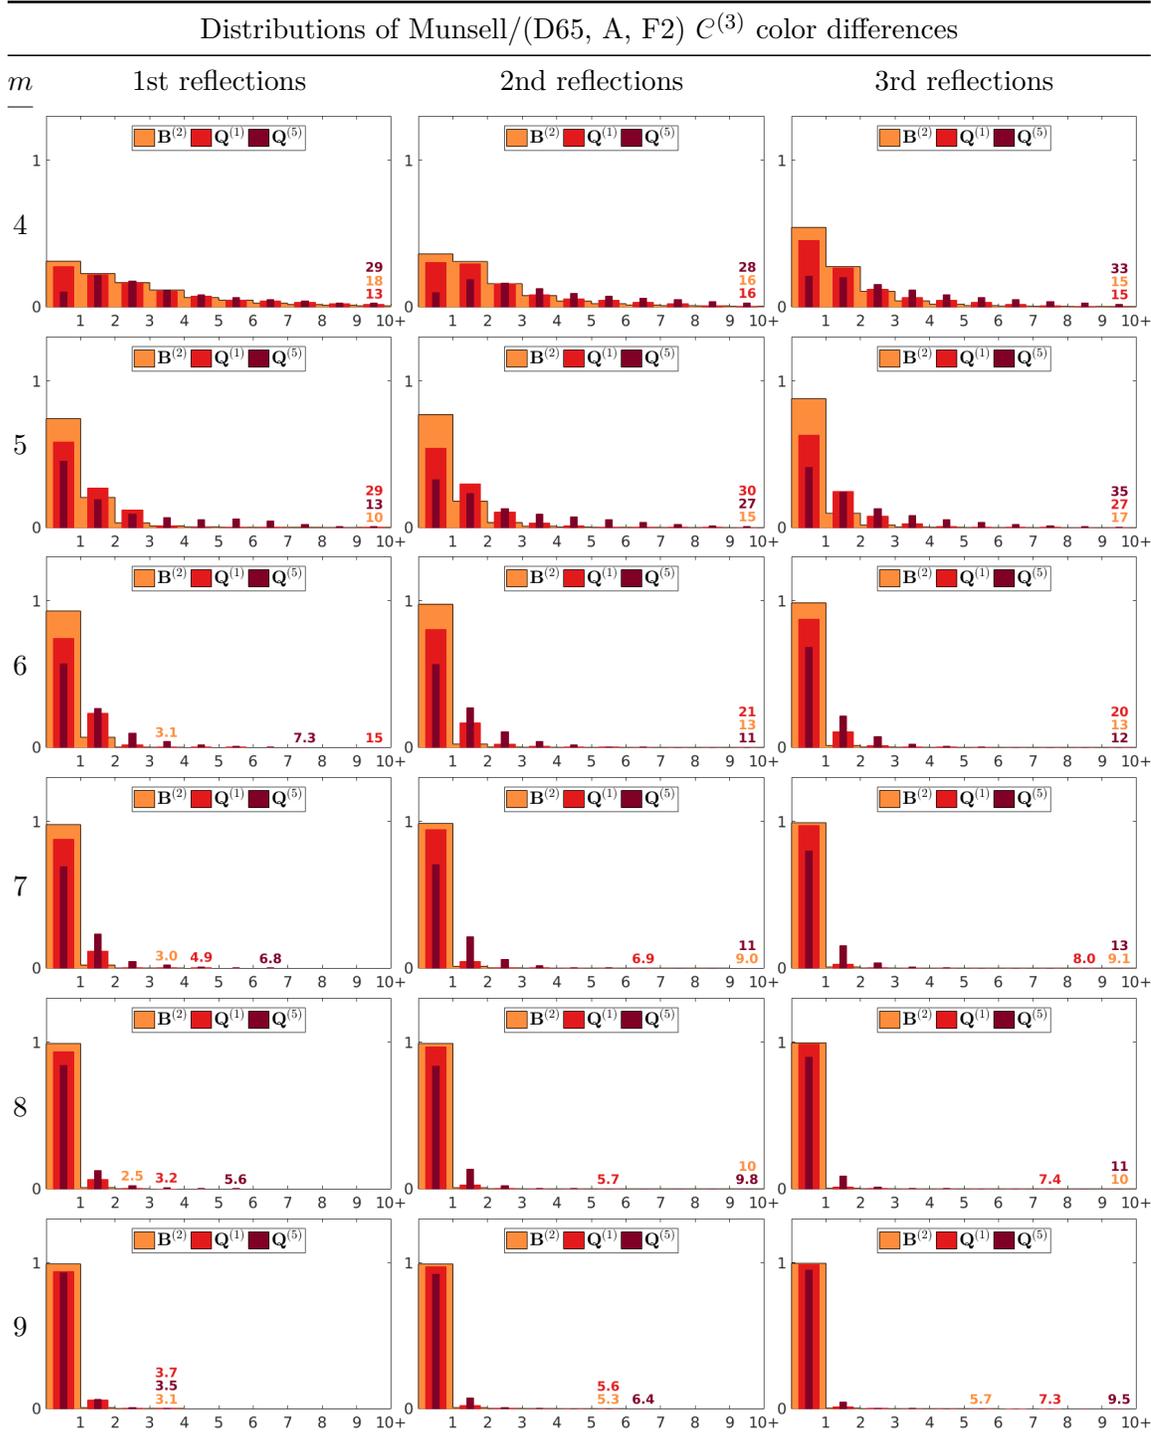


Figure A20: Distributions of ΔE_{00} color differences (as fractions of the totals) incurred when computing Munsell/(D65, A, F2) $C^{(3)}$ color signals by multiplying coefficients of $Q^{(1)}$ or $Q^{(5)}$, for dimensionalities $m \equiv 4, \dots, 9$ and first, second and third reflections. For comparison, differences incurred by projection onto $B^{(2)}$ are also shown. Maximum values, rounded to two significant figures, are shown in corresponding colors.

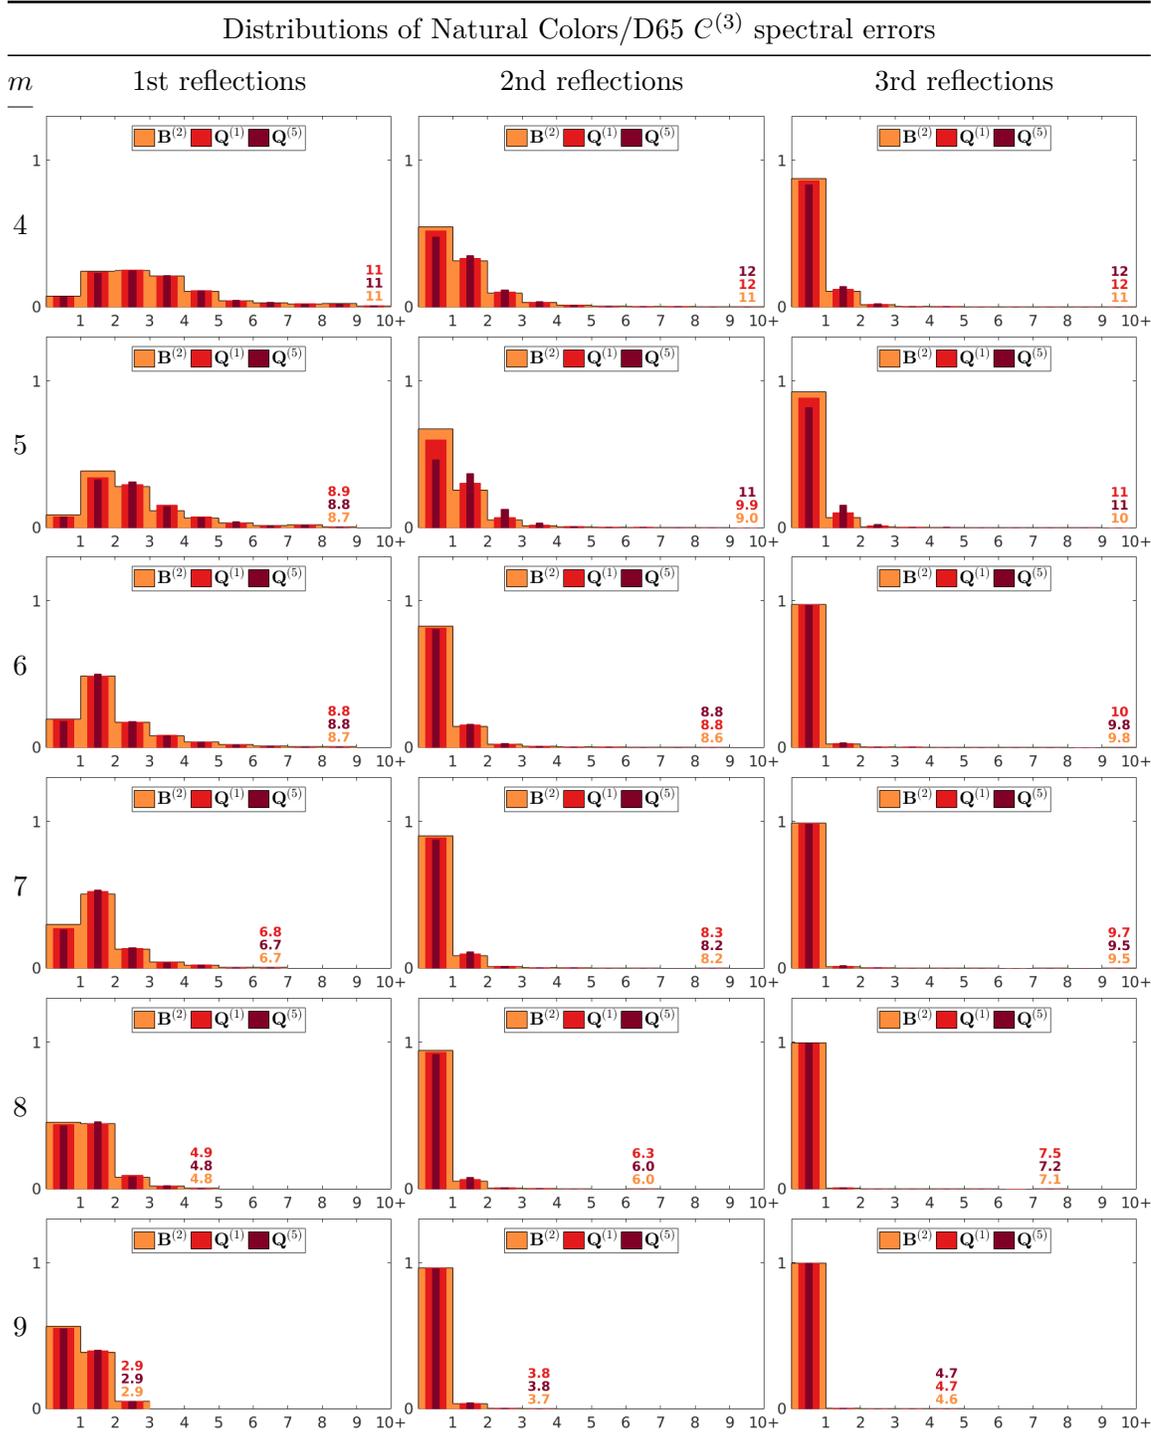


Figure A21: Distributions of NRMS spectral errors (as fractions of the totals) incurred when computing Natural Colors/D65 $C^{(3)}$ color signals by multiplying coefficients of $Q^{(1)}$ or $Q^{(5)}$, for dimensionalities $m \equiv 4, \dots, 9$ and first, second and third reflections. For comparison, errors incurred by projection onto $B^{(2)}$ are also shown. Maximum values, rounded to two significant figures, are shown in corresponding colors.

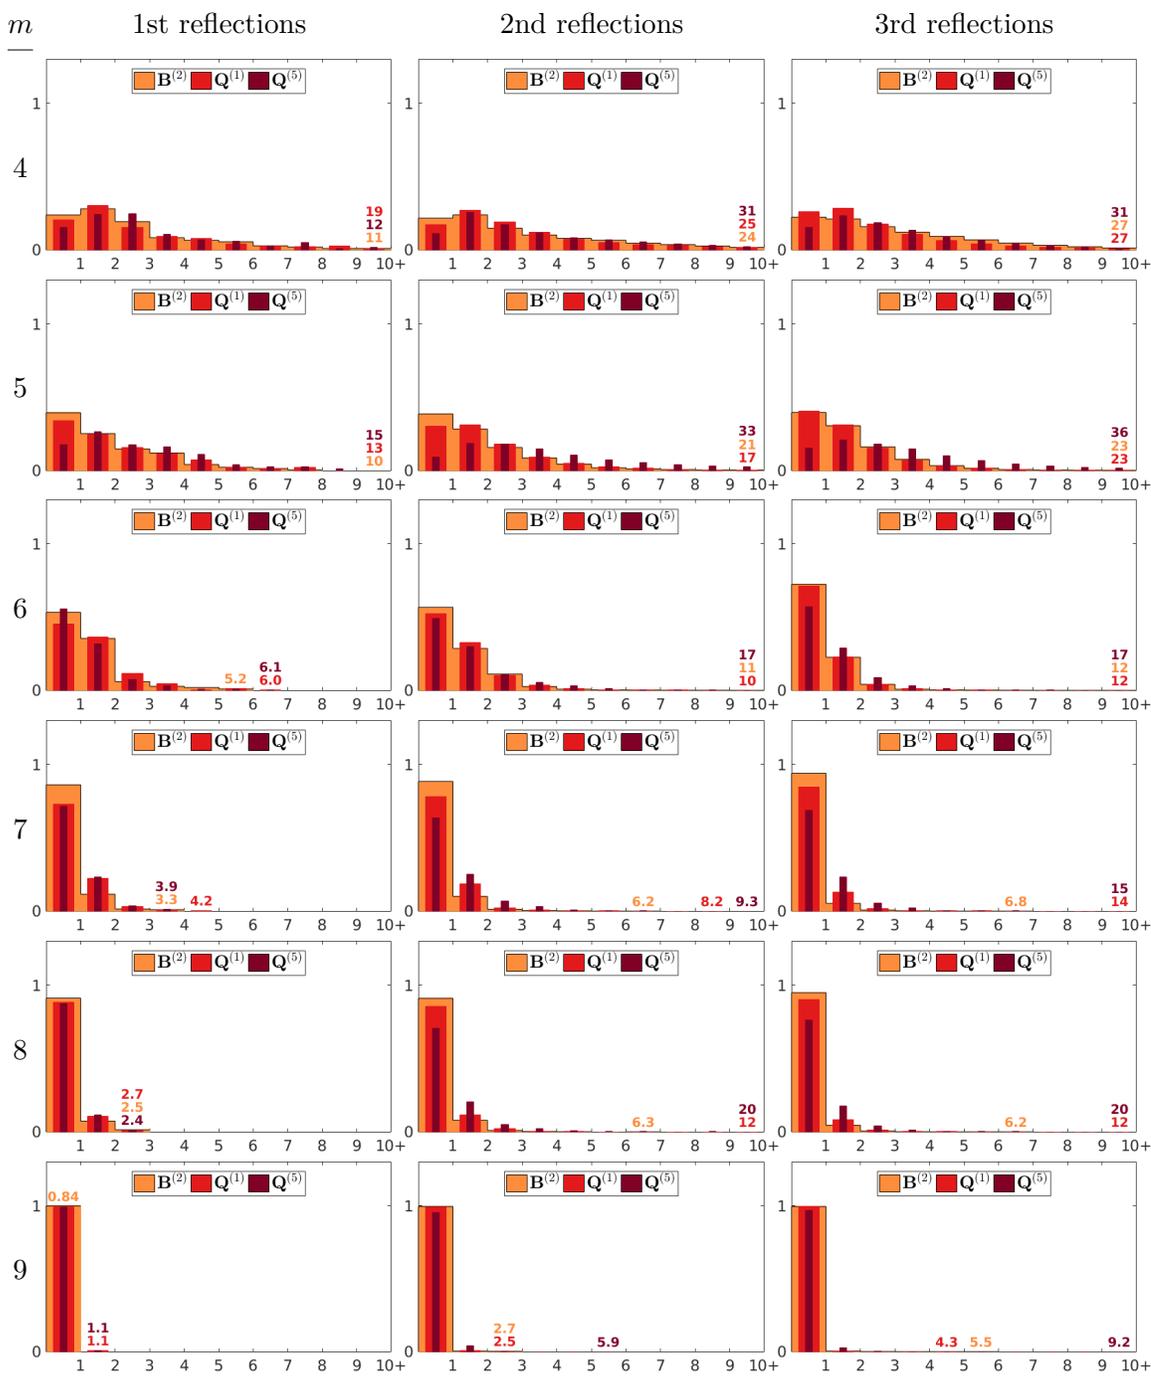
Distributions of Natural Colors/D65 $\mathcal{C}^{(3)}$ color differences

Figure A22: Distributions of ΔE_{00} color differences (as fractions of the totals) incurred when computing Natural Colors/D65 $\mathcal{C}^{(3)}$ color signals by multiplying coefficients of $\mathbf{Q}^{(1)}$ or $\mathbf{Q}^{(5)}$, for dimensionalities $m \equiv 4, \dots, 9$ and first, second and third reflections. For comparison, differences incurred by projection onto $\mathbf{B}^{(2)}$ are also shown. Maximum values, rounded to two significant figures, are shown in corresponding colors.

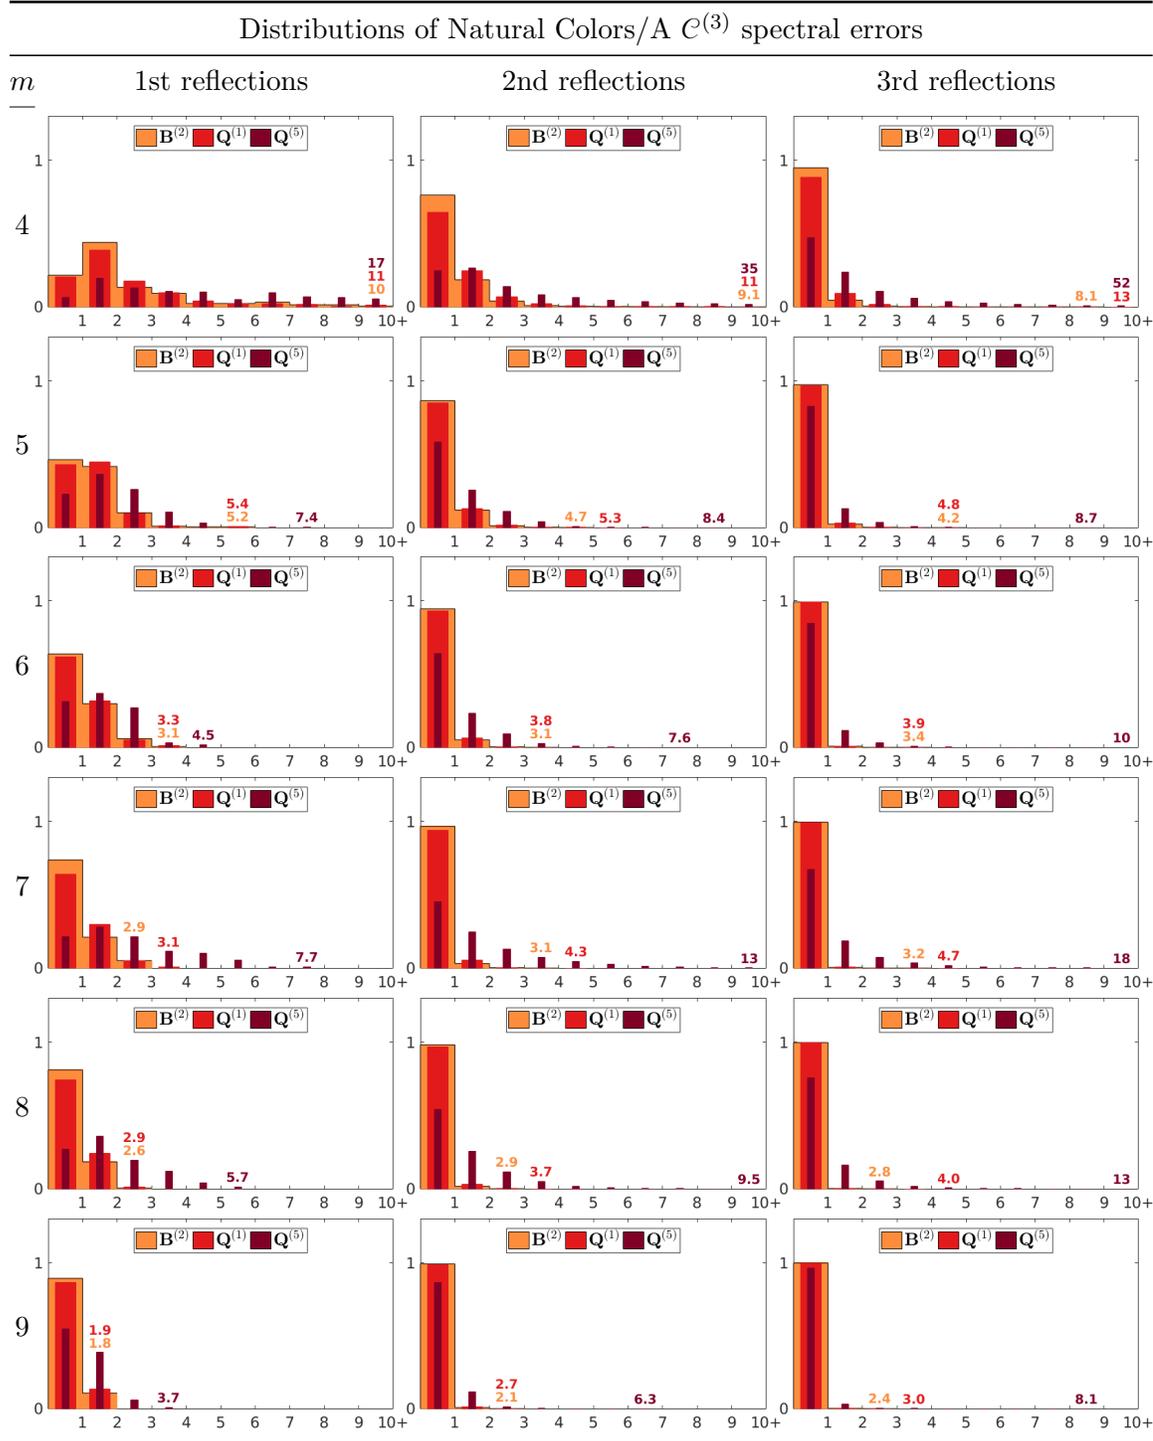


Figure A23: Distributions of NRMS spectral errors (as fractions of the totals) incurred when computing Natural Colors/A $C^{(3)}$ color signals by multiplying coefficients of $Q^{(1)}$ or $Q^{(5)}$, for dimensionalities $m \equiv 4, \dots, 9$ and first, second and third reflections. For comparison, errors incurred by projection onto $B^{(2)}$ are also shown. Maximum values, rounded to two significant figures, are shown in corresponding colors.

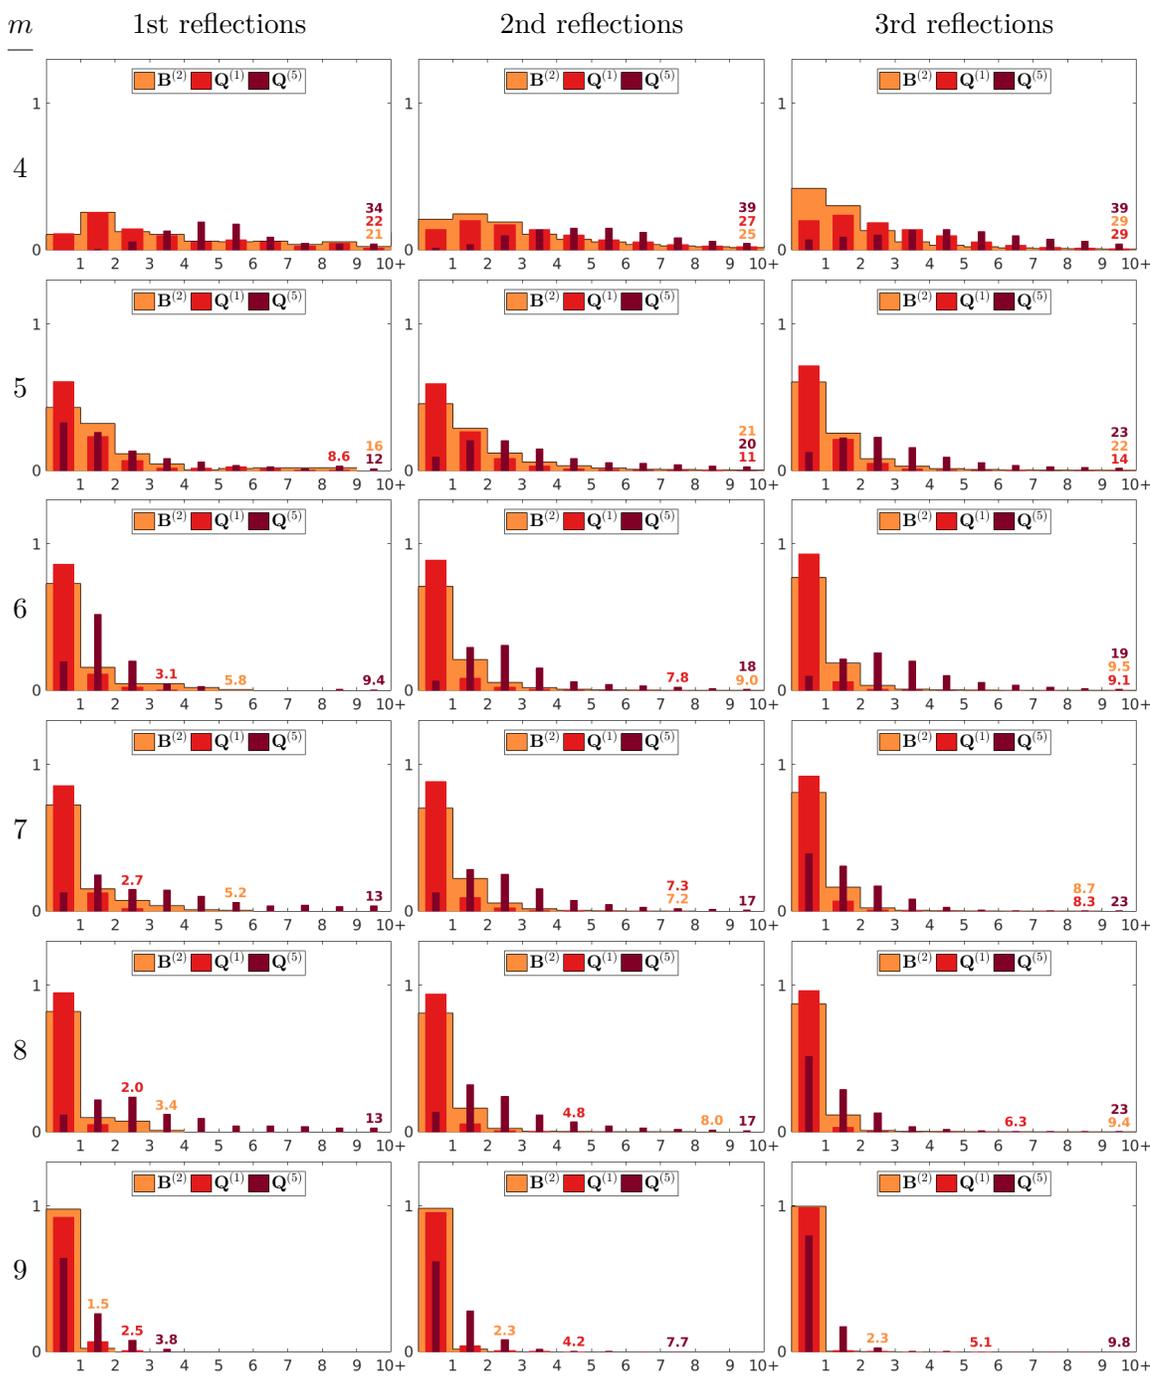
Distributions of Natural Colors/A $C^{(3)}$ color differences

Figure A24: Distributions of ΔE_{00} color differences (as fractions of the totals) incurred when computing Natural Colors/A $C^{(3)}$ color signals by multiplying coefficients of $Q^{(1)}$ or $Q^{(5)}$, for dimensionalities $m \equiv 4, \dots, 9$ and first, second and third reflections. For comparison, differences incurred by projection onto $B^{(2)}$ are also shown. Maximum values, rounded to two significant figures, are shown in corresponding colors.

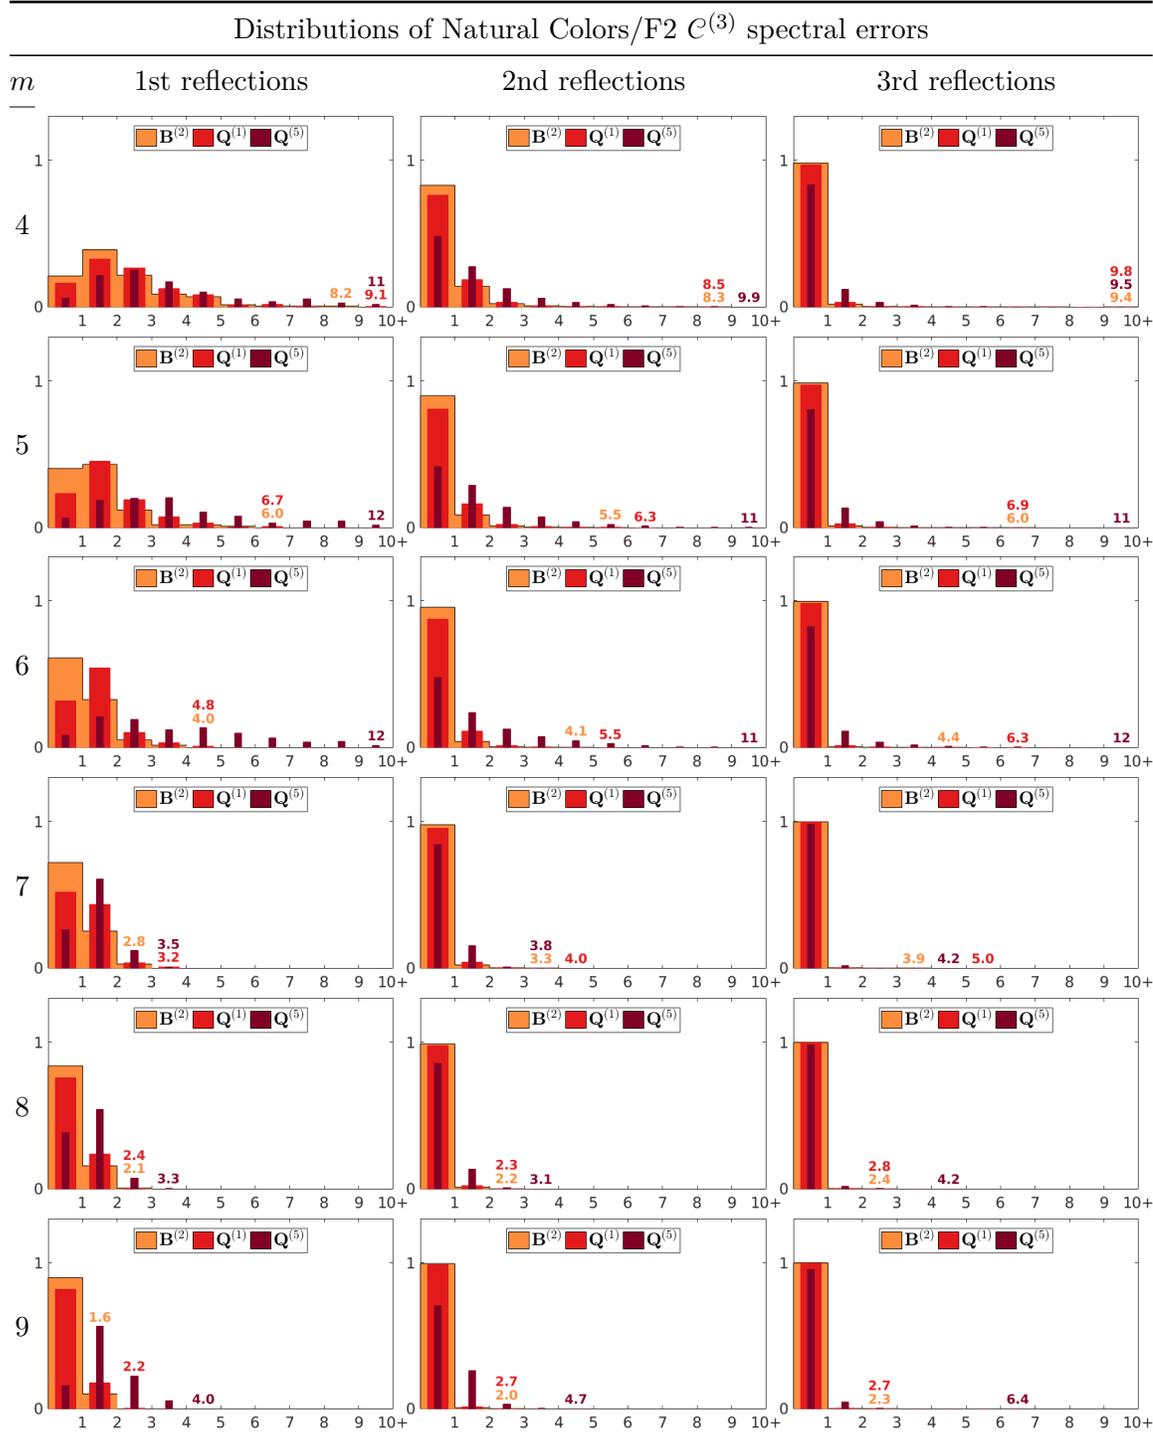


Figure A25: Distributions of NRMS spectral errors (as fractions of the totals) incurred when computing Natural Colors/F2 $C^{(3)}$ color signals by multiplying coefficients of $Q^{(1)}$ or $Q^{(5)}$, for dimensionalities $m \equiv 4, \dots, 9$ and first, second and third reflections. For comparison, errors incurred by projection onto $B^{(2)}$ are also shown. Maximum values, rounded to two significant figures, are shown in corresponding colors.

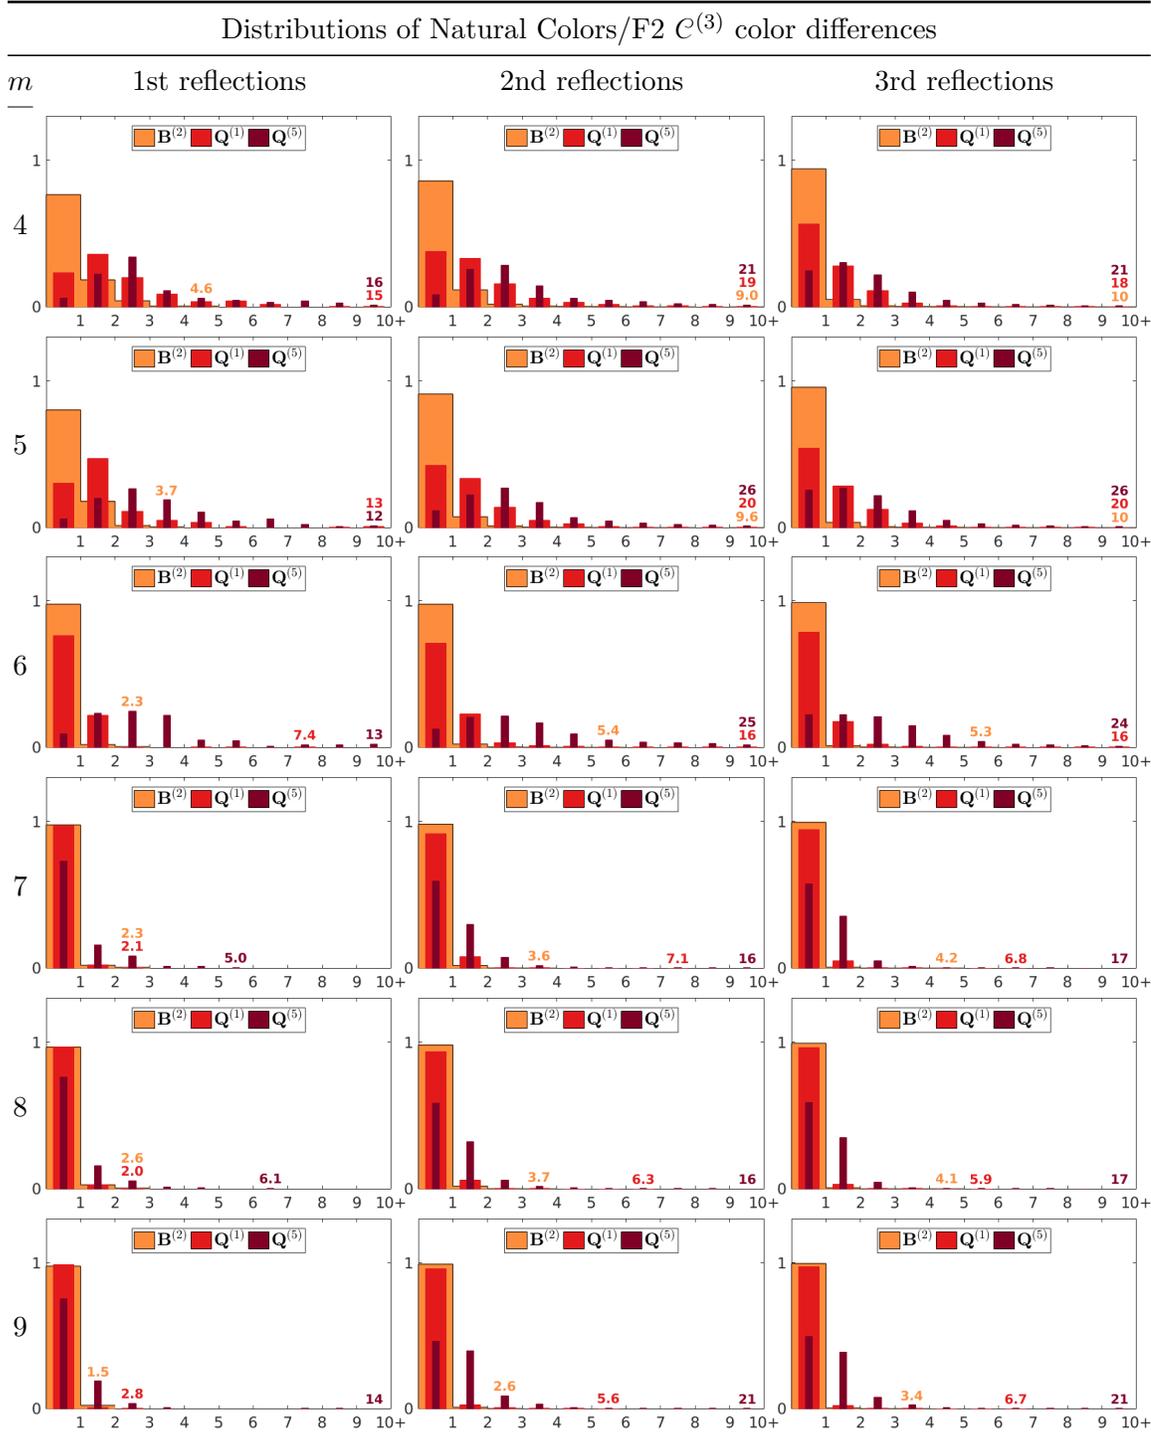


Figure A26: Distributions of ΔE_{00} color differences (as fractions of the totals) incurred when computing Natural Colors/F2 $C^{(3)}$ color signals by multiplying coefficients of $Q^{(1)}$ or $Q^{(5)}$, for dimensionalities $m \equiv 4, \dots, 9$ and first, second and third reflections. For comparison, differences incurred by projection onto $B^{(2)}$ are also shown. Maximum values, rounded to two significant figures, are shown in corresponding colors.

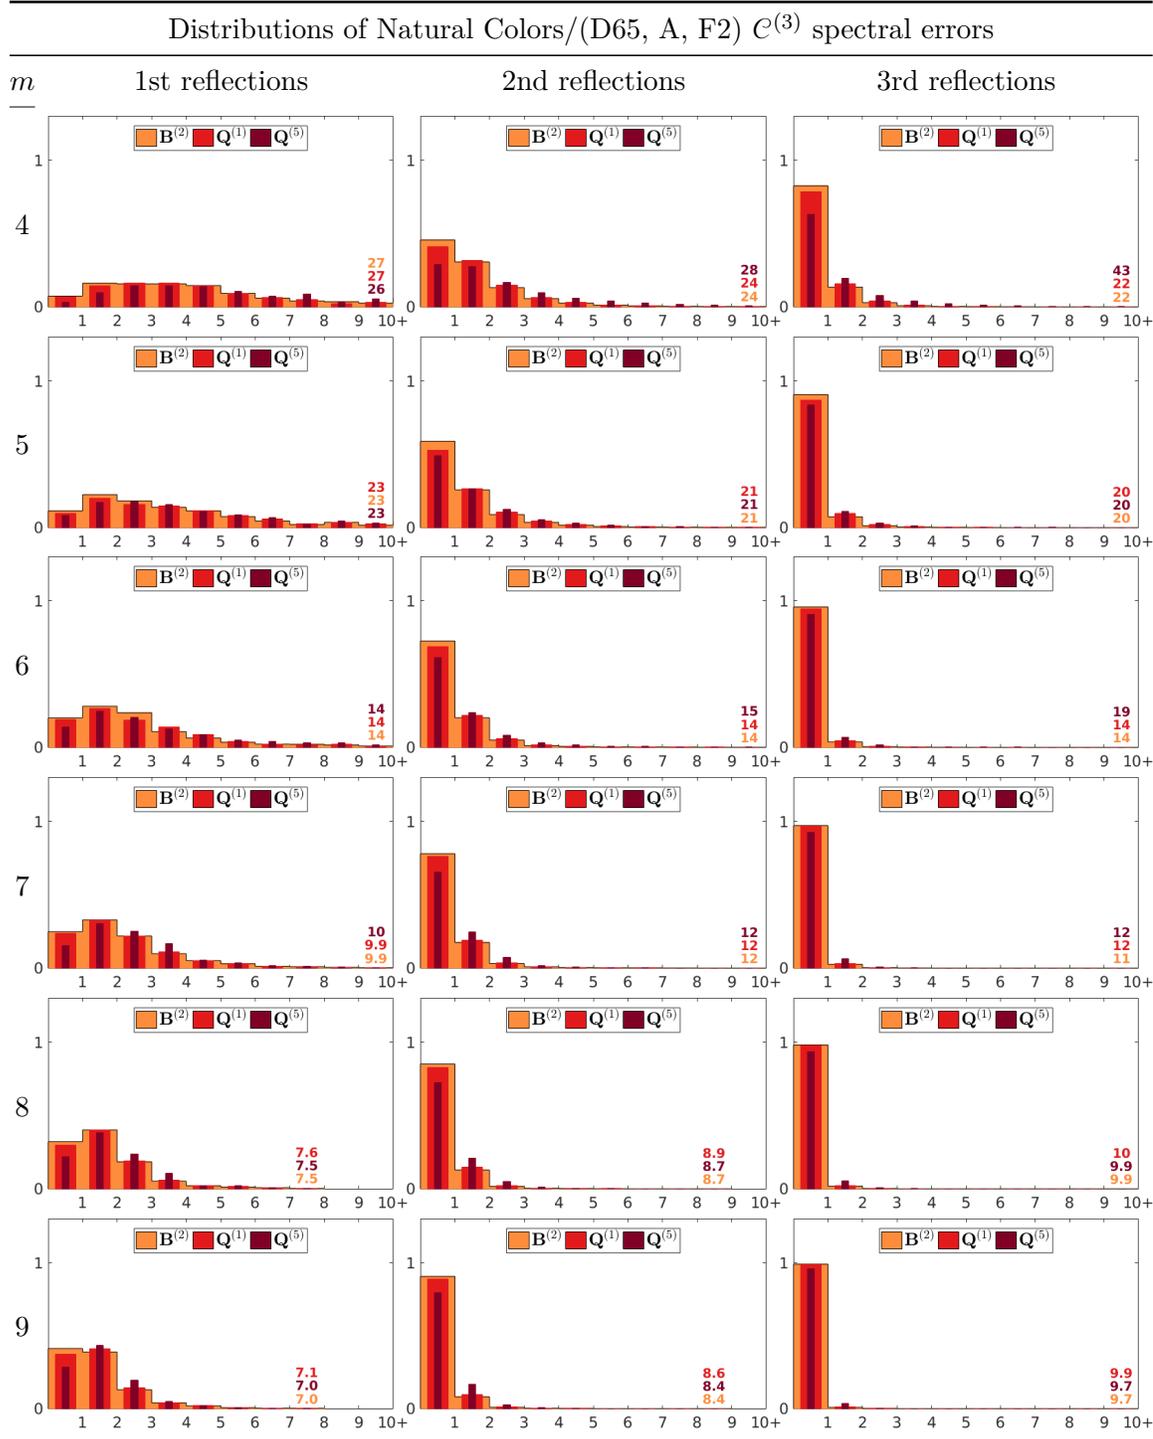


Figure A27: Distributions of NRMS spectral errors (as fractions of the totals) incurred when computing Natural Colors/(D65, A, F2) $\mathcal{C}^{(3)}$ color signals by multiplying coefficients of $\mathbf{Q}^{(1)}$ or $\mathbf{Q}^{(5)}$, for dimensionalities $m \equiv 4, \dots, 9$ and first, second and third reflections. For comparison, errors incurred by projection onto $\mathbf{B}^{(2)}$ are also shown. Maximum values, rounded to two significant figures, are shown in corresponding colors.

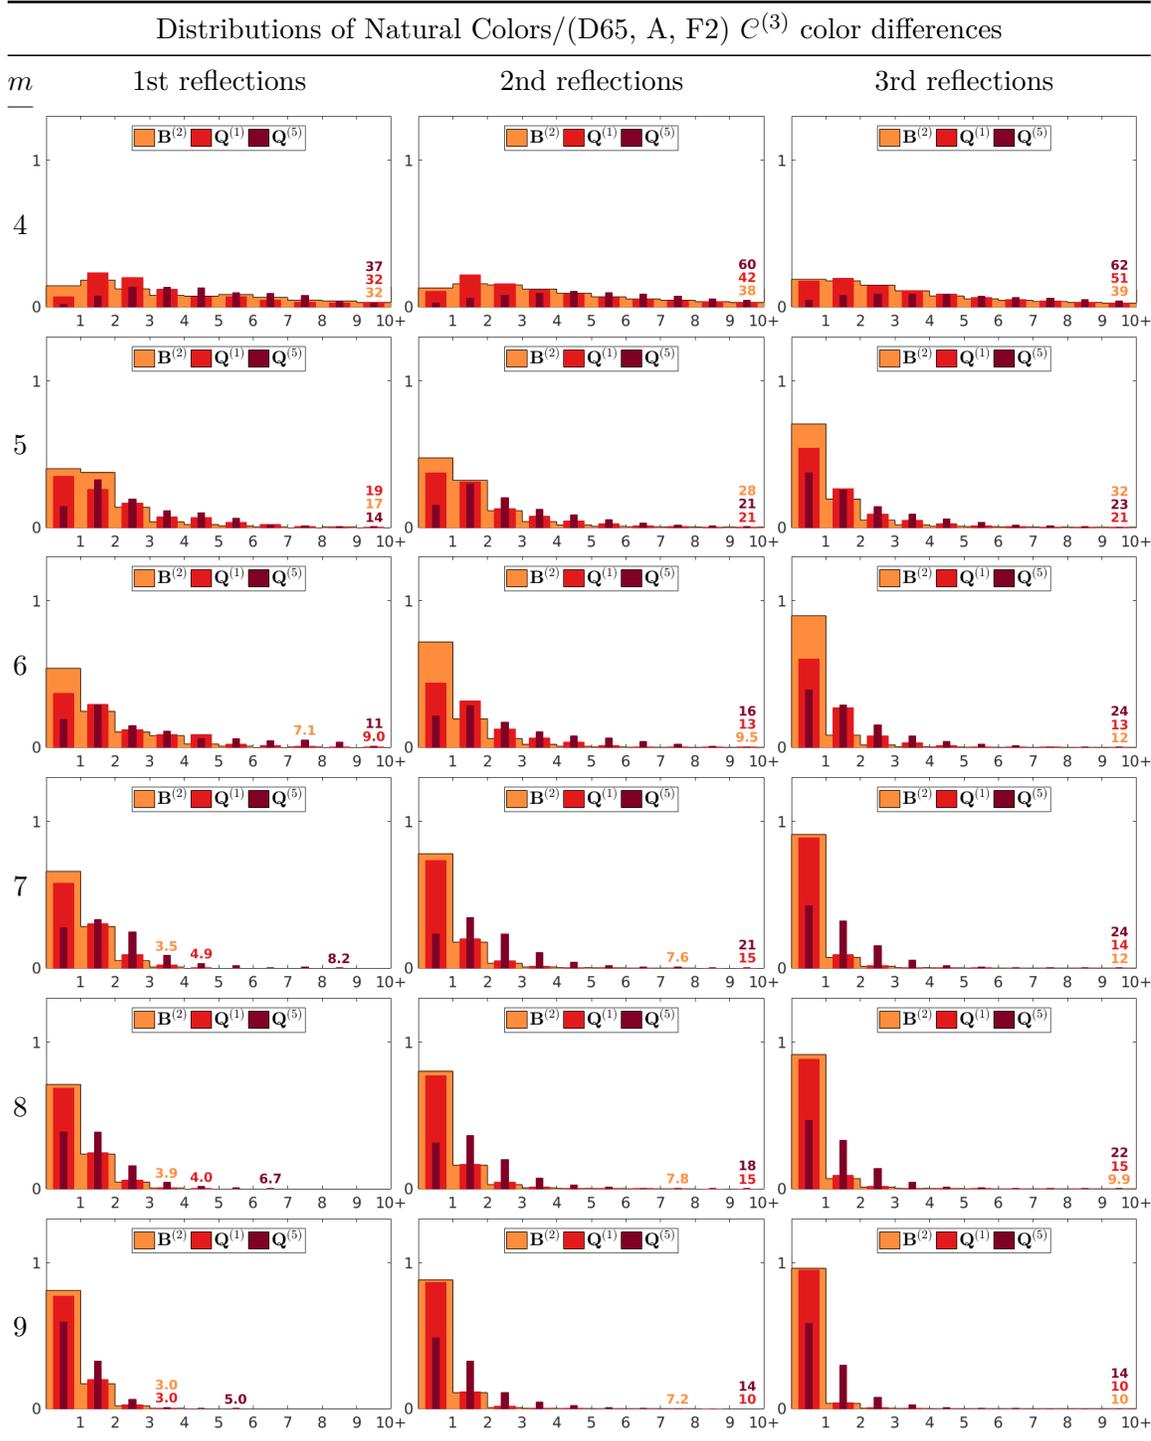


Figure A28: Distributions of ΔE_{00} color differences (as fractions of the totals) incurred when computing Natural Colors/(D65, A, F2) $C^{(3)}$ color signals by multiplying coefficients of $Q^{(1)}$ or $Q^{(5)}$, for dimensionalities $m \equiv 4, \dots, 9$ and first, second and third reflections. For comparison, differences incurred by projection onto $B^{(2)}$ are also shown. Maximum values, rounded to two significant figures, are shown in corresponding colors.

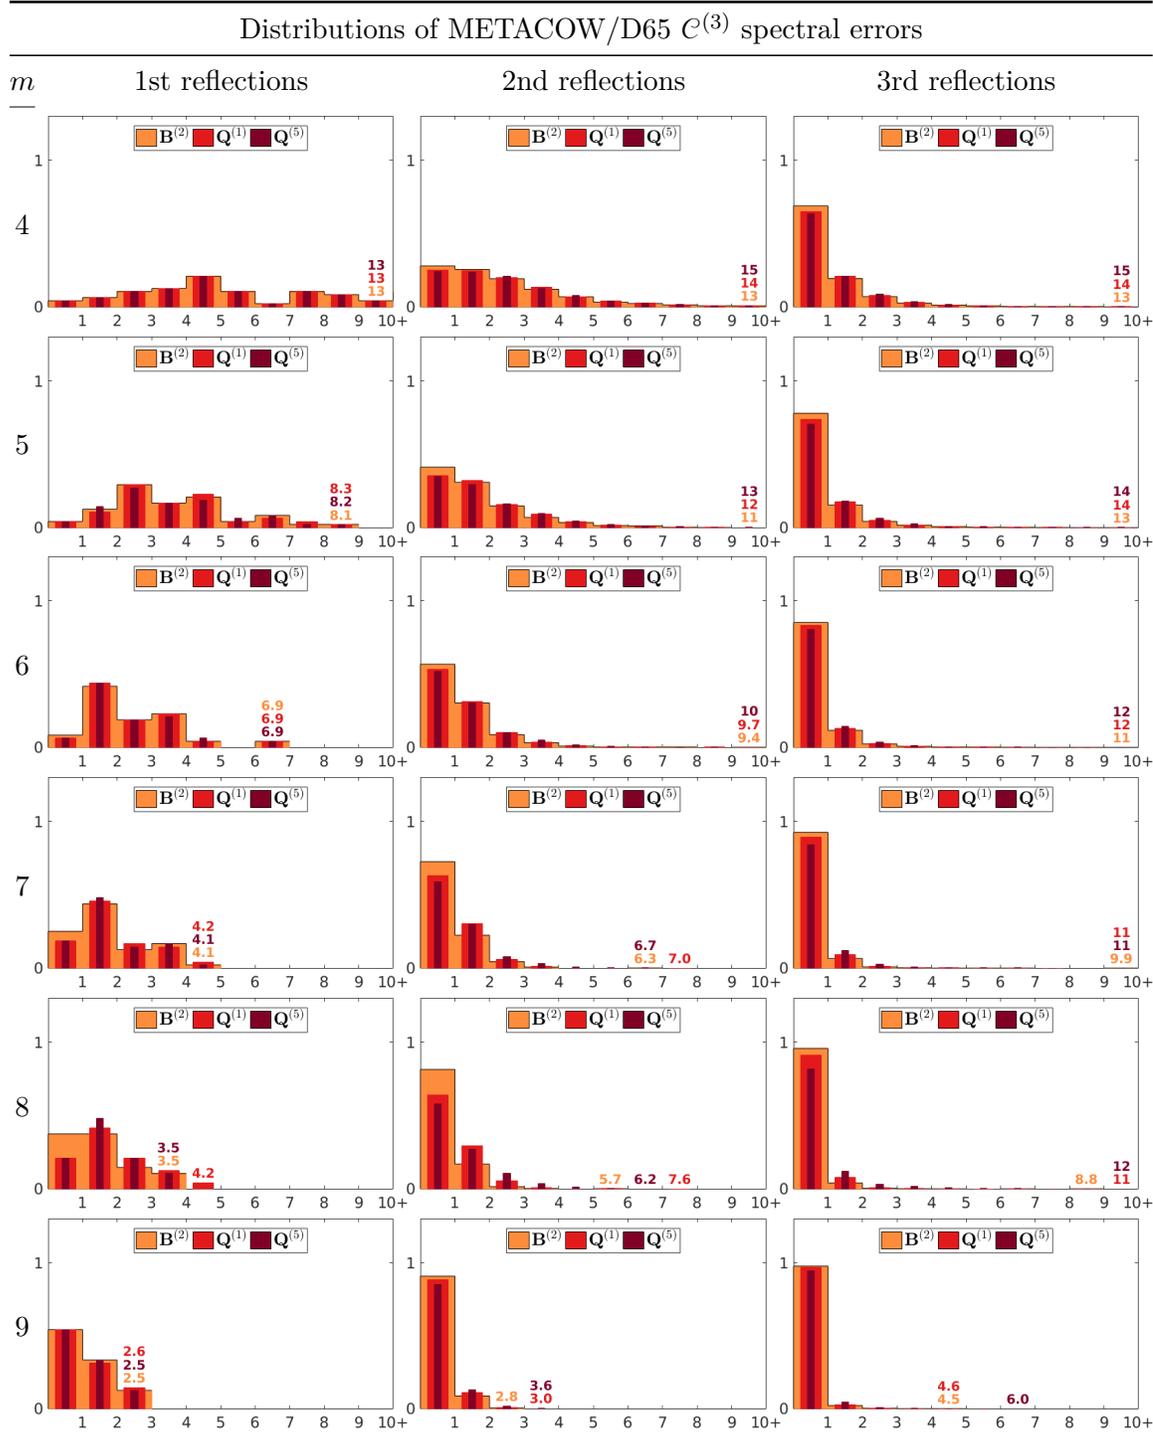


Figure A29: Distributions of NRMS spectral errors (as fractions of the totals) incurred when computing METACOW/D65 $C^{(3)}$ color signals by multiplying coefficients of $\mathbf{Q}^{(1)}$ or $\mathbf{Q}^{(5)}$, for dimensionalities $m \equiv 4, \dots, 9$ and first, second and third reflections. For comparison, errors incurred by projection onto $\mathbf{B}^{(2)}$ are also shown. Maximum values, rounded to two significant figures, are shown in corresponding colors.

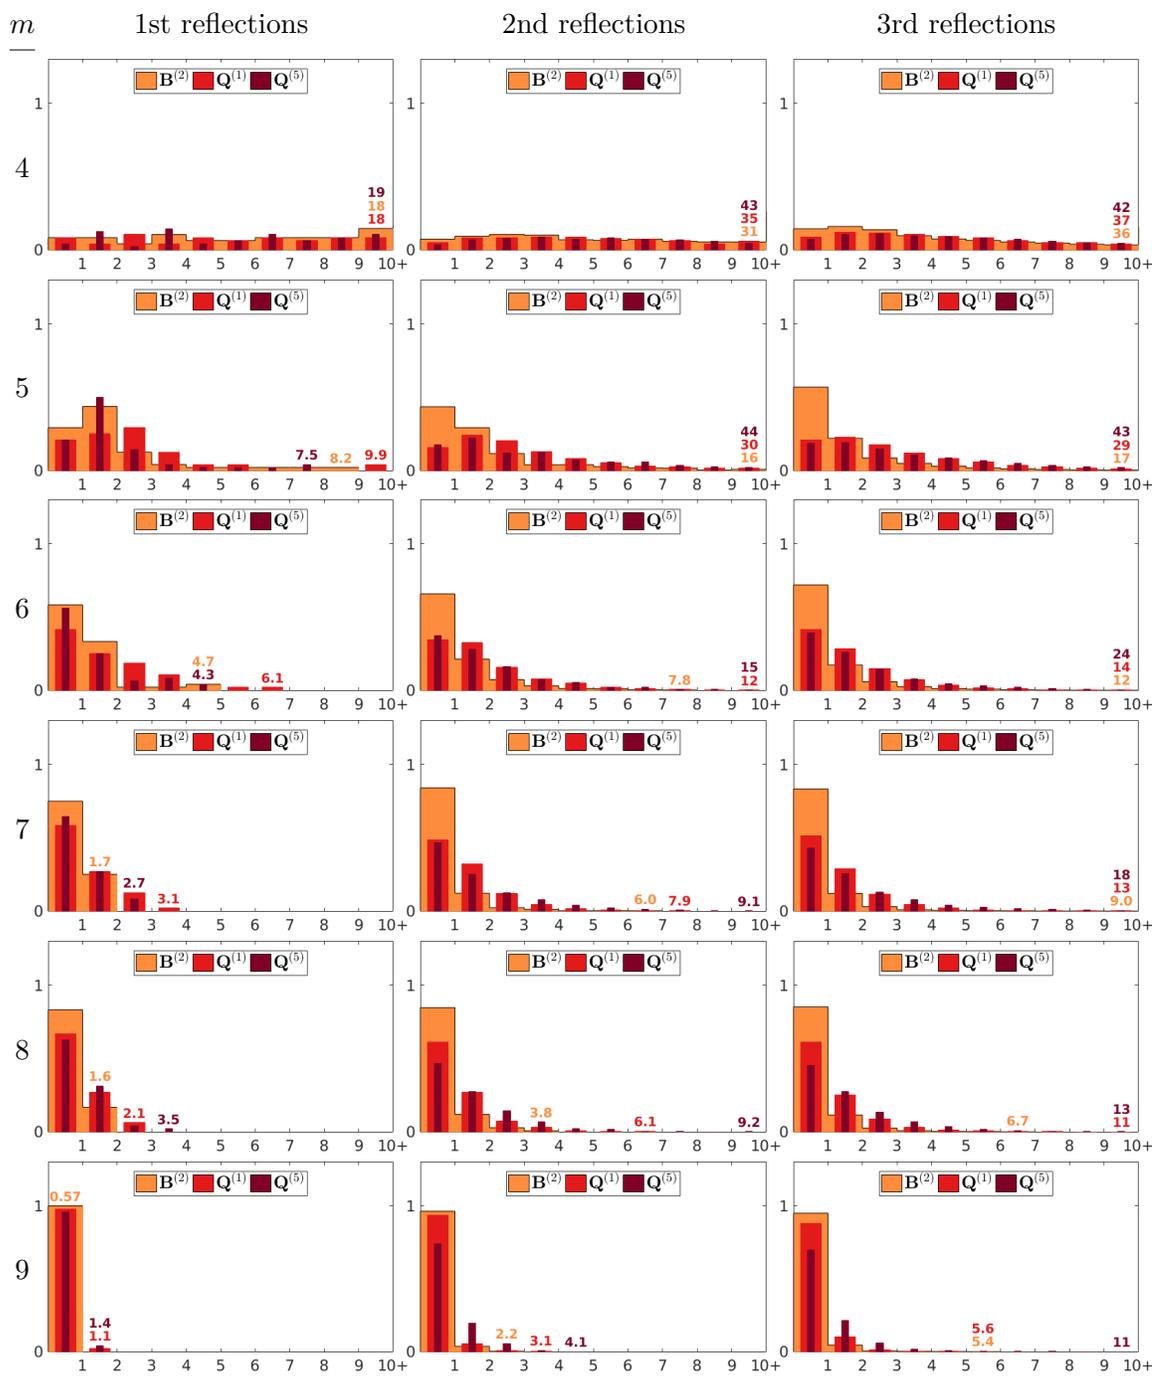
Distributions of METACOW/D65 $c^{(3)}$ color differences

Figure A30: Distributions of ΔE_{00} color differences (as fractions of the totals) incurred when computing METACOW/D65 $c^{(3)}$ color signals by multiplying coefficients of $Q^{(1)}$ or $Q^{(5)}$, for dimensionalities $m \equiv 4, \dots, 9$ and first, second and third reflections. For comparison, differences incurred by projection onto $B^{(2)}$ are also shown. Maximum values, rounded to two significant figures, are shown in corresponding colors.

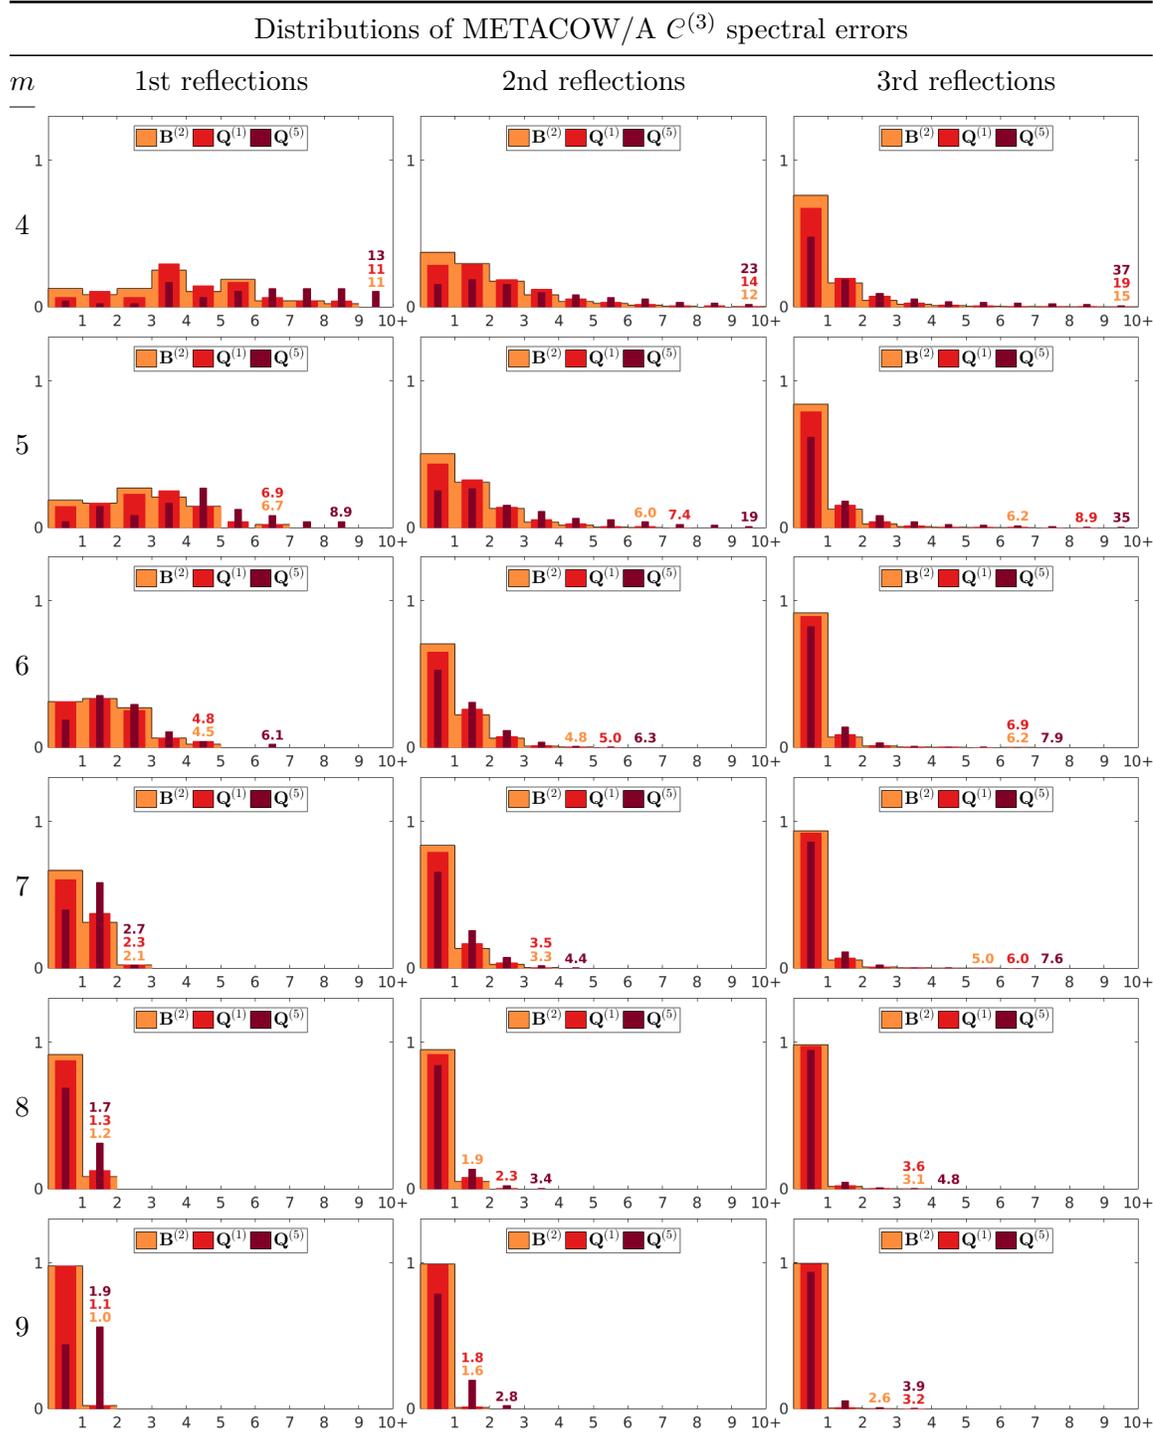


Figure A31: Distributions of NRMS spectral errors (as fractions of the totals) incurred when computing METACOW/A $c^{(3)}$ color signals by multiplying coefficients of $Q^{(1)}$ or $Q^{(5)}$, for dimensionalities $m \equiv 4, \dots, 9$ and first, second and third reflections. For comparison, errors incurred by projection onto $B^{(2)}$ are also shown. Maximum values, rounded to two significant figures, are shown in corresponding colors.

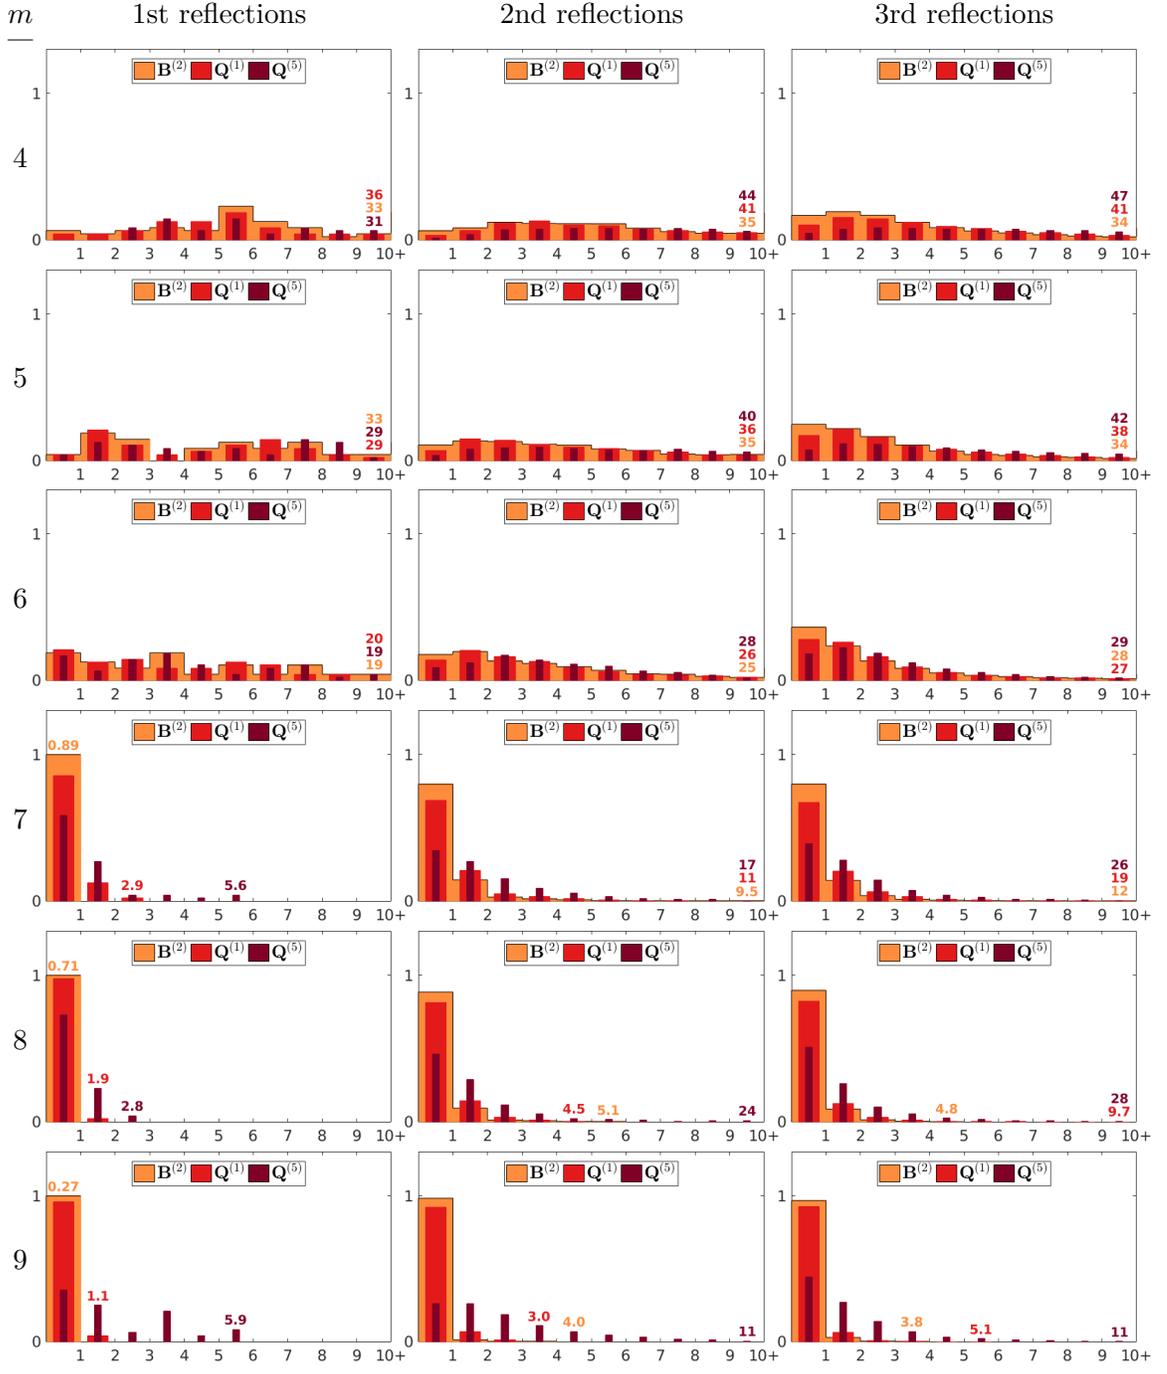
Distributions of METACOW/A $C^{(3)}$ color differences

Figure A32: Distributions of ΔE_{00} color differences (as fractions of the totals) incurred when computing METACOW/A $C^{(3)}$ color signals by multiplying coefficients of $\mathbf{Q}^{(1)}$ or $\mathbf{Q}^{(5)}$, for dimensionalities $m \equiv 4, \dots, 9$ and first, second and third reflections. For comparison, differences incurred by projection onto $\mathbf{B}^{(2)}$ are also shown. Maximum values, rounded to two significant figures, are shown in corresponding colors.

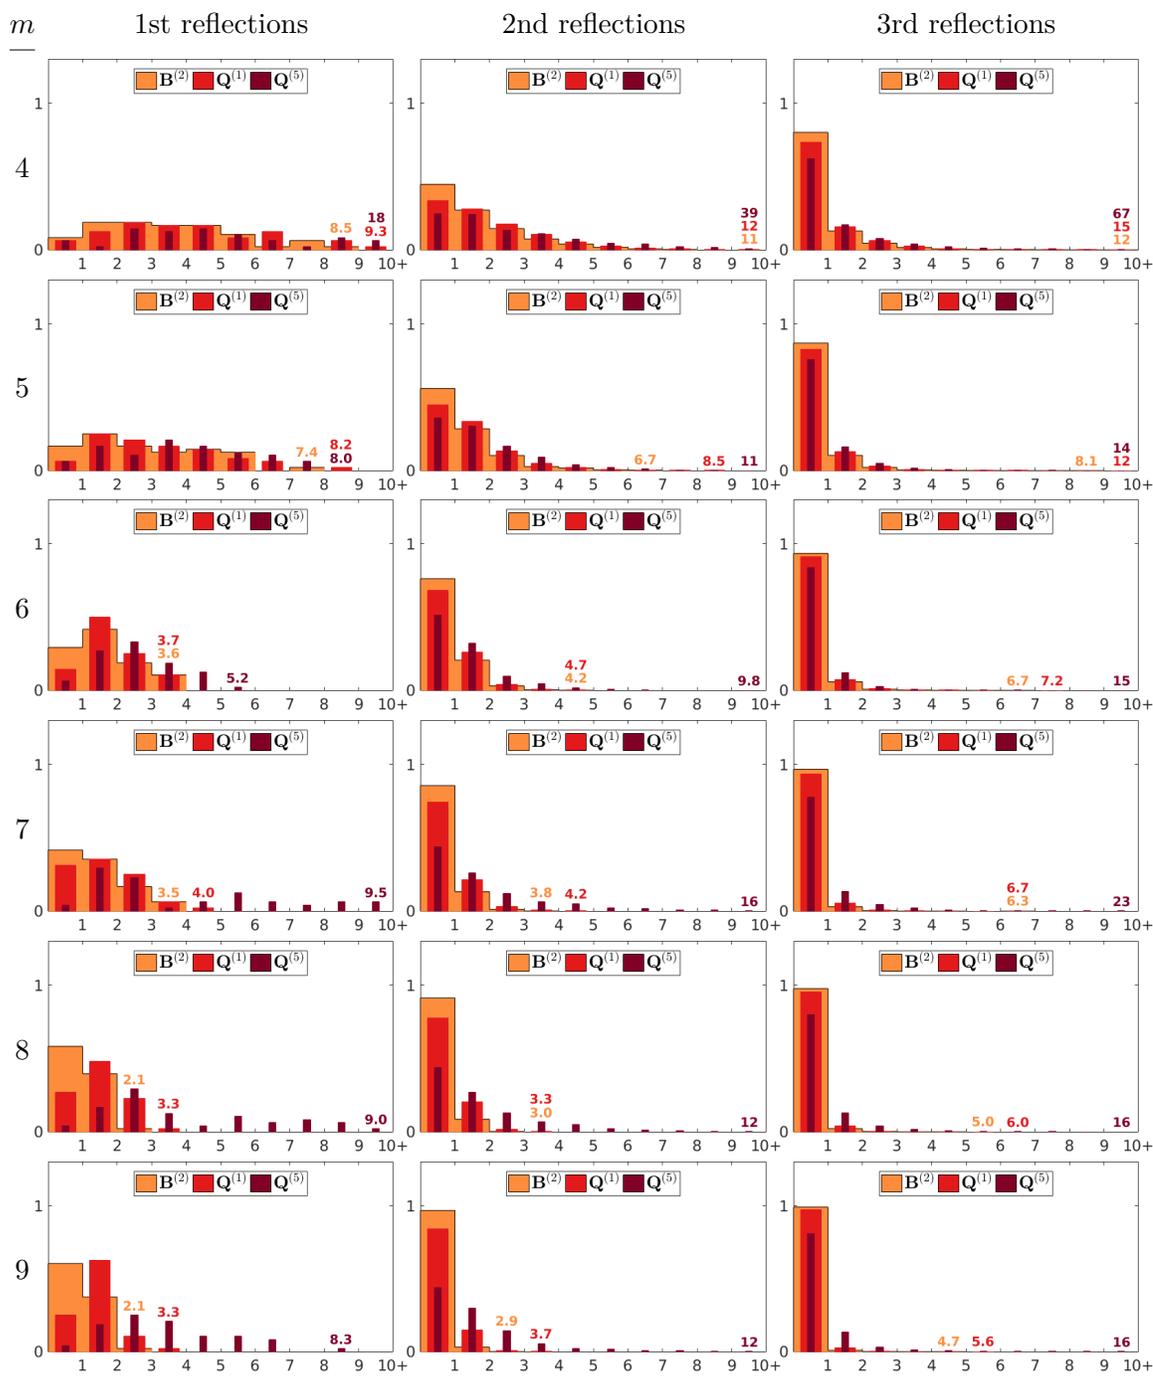
Distributions of METACOW/F2 $\mathcal{C}^{(3)}$ spectral errors

Figure A33: Distributions of NRMS spectral errors (as fractions of the totals) incurred when computing METACOW/F2 $\mathcal{C}^{(3)}$ color signals by multiplying coefficients of $\mathbf{Q}^{(1)}$ or $\mathbf{Q}^{(5)}$, for dimensionalities $m \equiv 4, \dots, 9$ and first, second and third reflections. For comparison, errors incurred by projection onto $\mathbf{B}^{(2)}$ are also shown. Maximum values, rounded to two significant figures, are shown in corresponding colors.

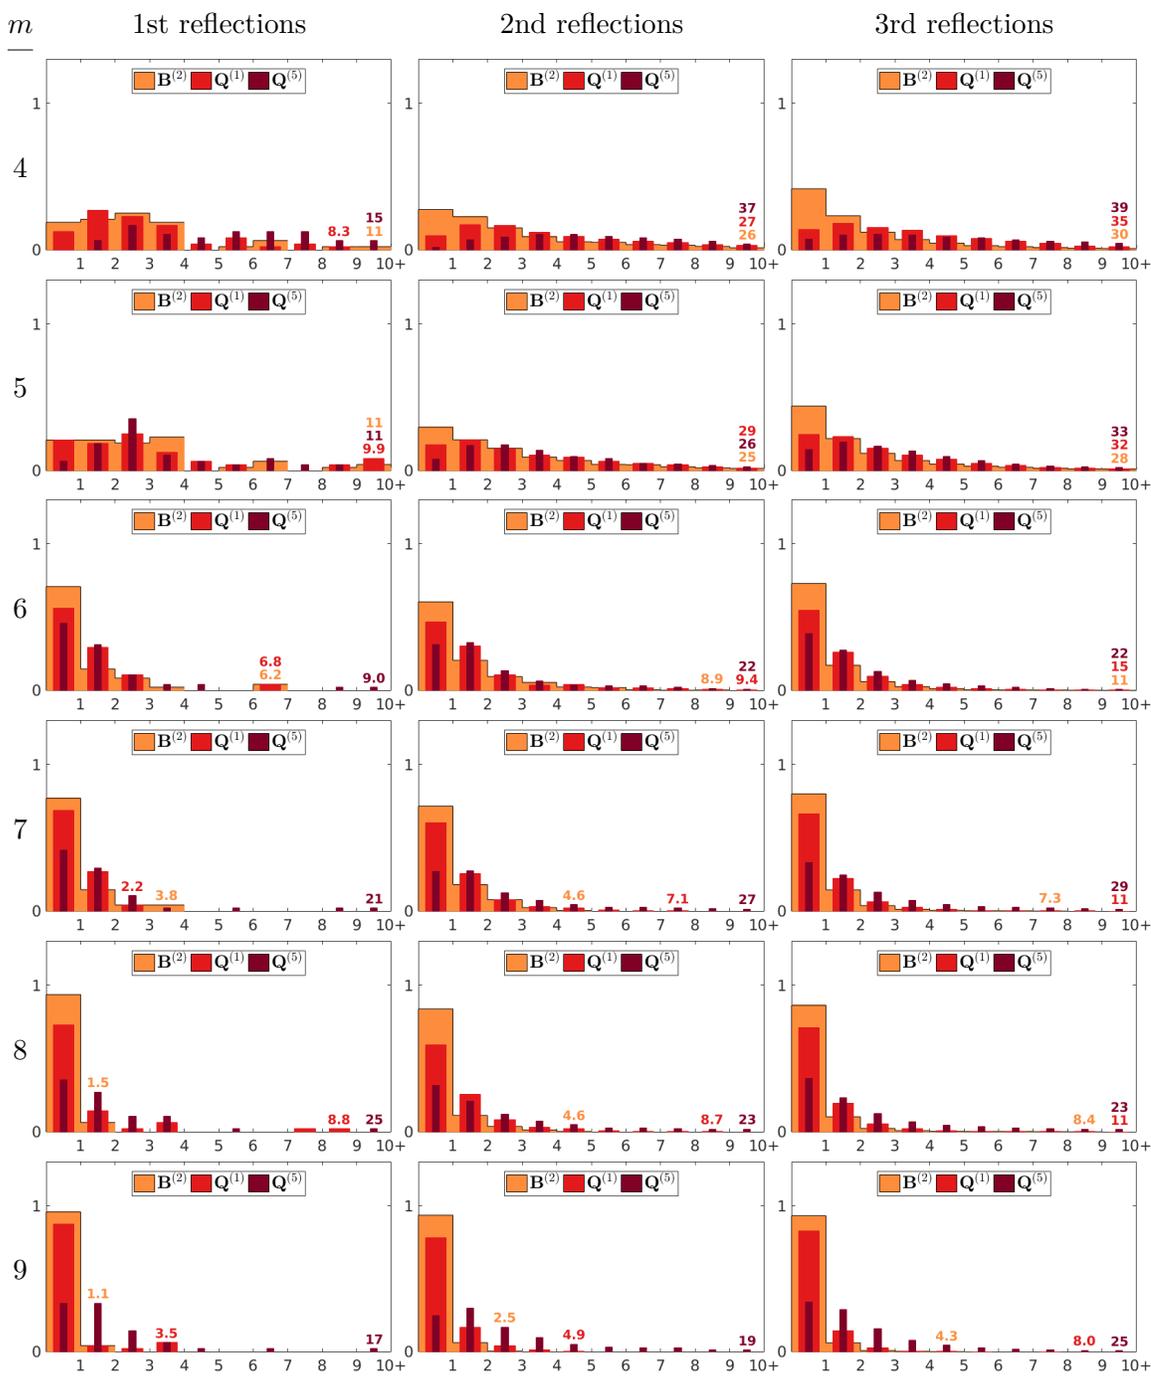
Distributions of METACOW/F2 $C^{(3)}$ color differences

Figure A34: Distributions of ΔE_{00} color differences (as fractions of the totals) incurred when computing METACOW/F2 $C^{(3)}$ color signals by multiplying coefficients of $Q^{(1)}$ or $Q^{(5)}$, for dimensionalities $m \equiv 4, \dots, 9$ and first, second and third reflections. For comparison, differences incurred by projection onto $B^{(2)}$ are also shown. Maximum values, rounded to two significant figures, are shown in corresponding colors.

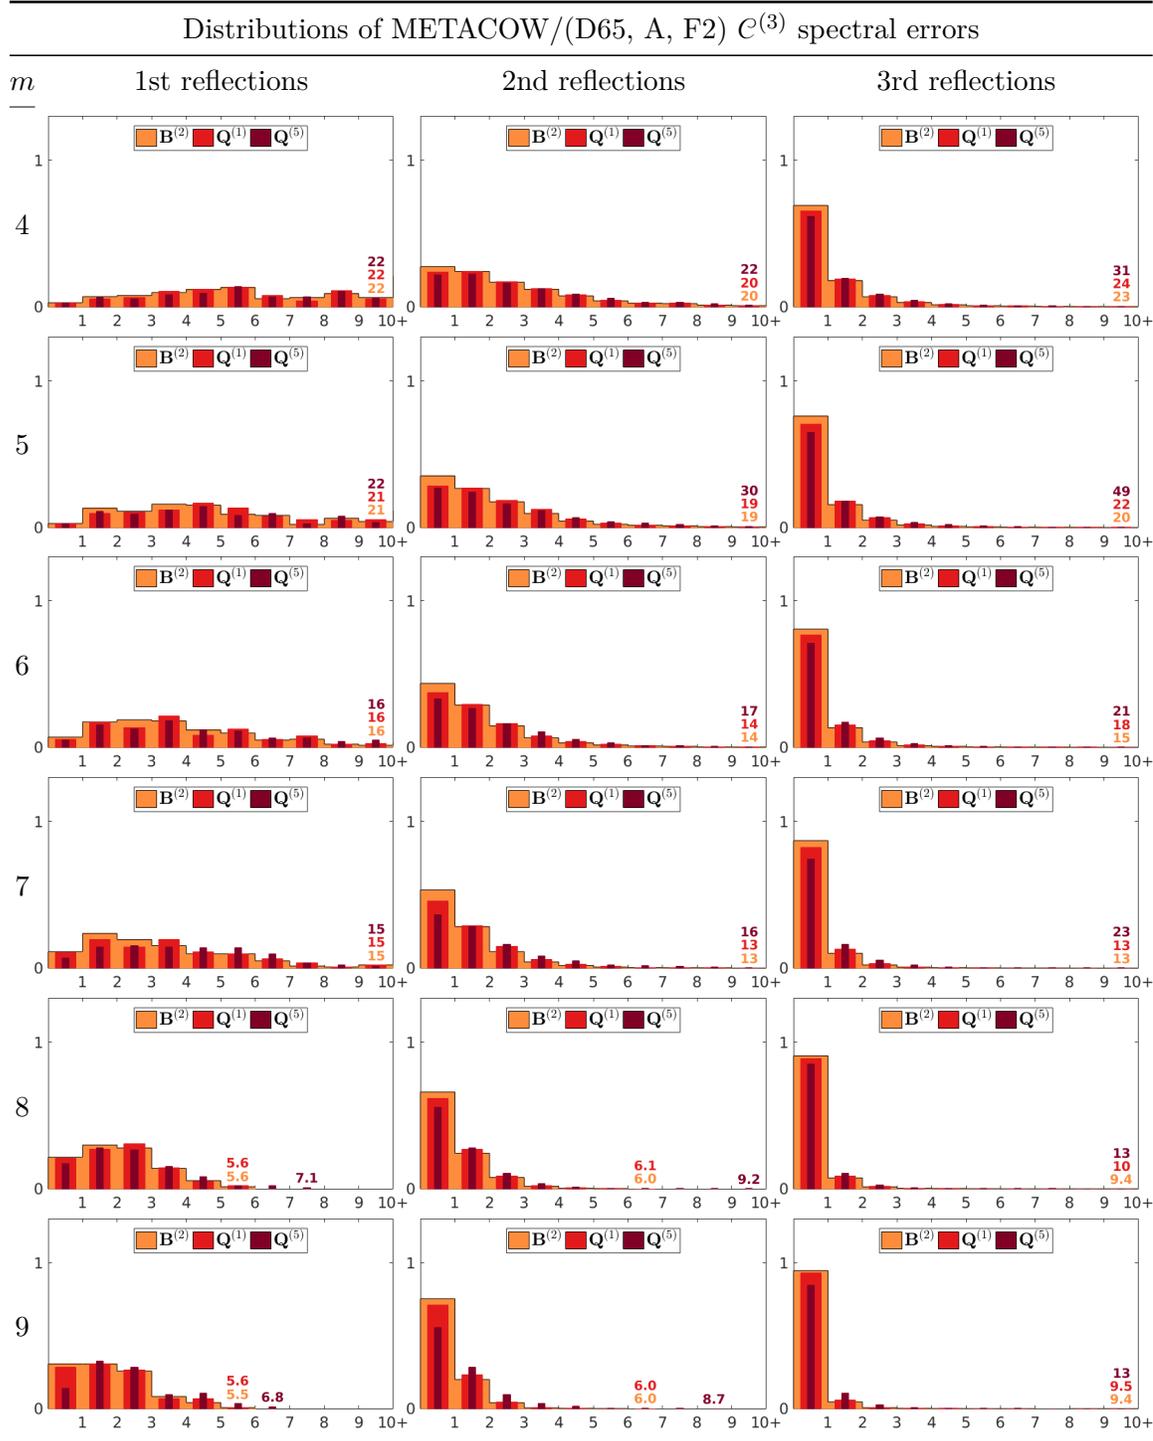


Figure A35: Distributions of NRMS spectral errors (as fractions of the totals) incurred when computing METACOW/(D65, A, F2) $C^{(3)}$ color signals by multiplying coefficients of $\mathbf{Q}^{(1)}$ or $\mathbf{Q}^{(5)}$, for dimensionalities $m \equiv 4, \dots, 9$ and first, second and third reflections. For comparison, errors incurred by projection onto $\mathbf{B}^{(2)}$ are also shown. Maximum values, rounded to two significant figures, are shown in corresponding colors.

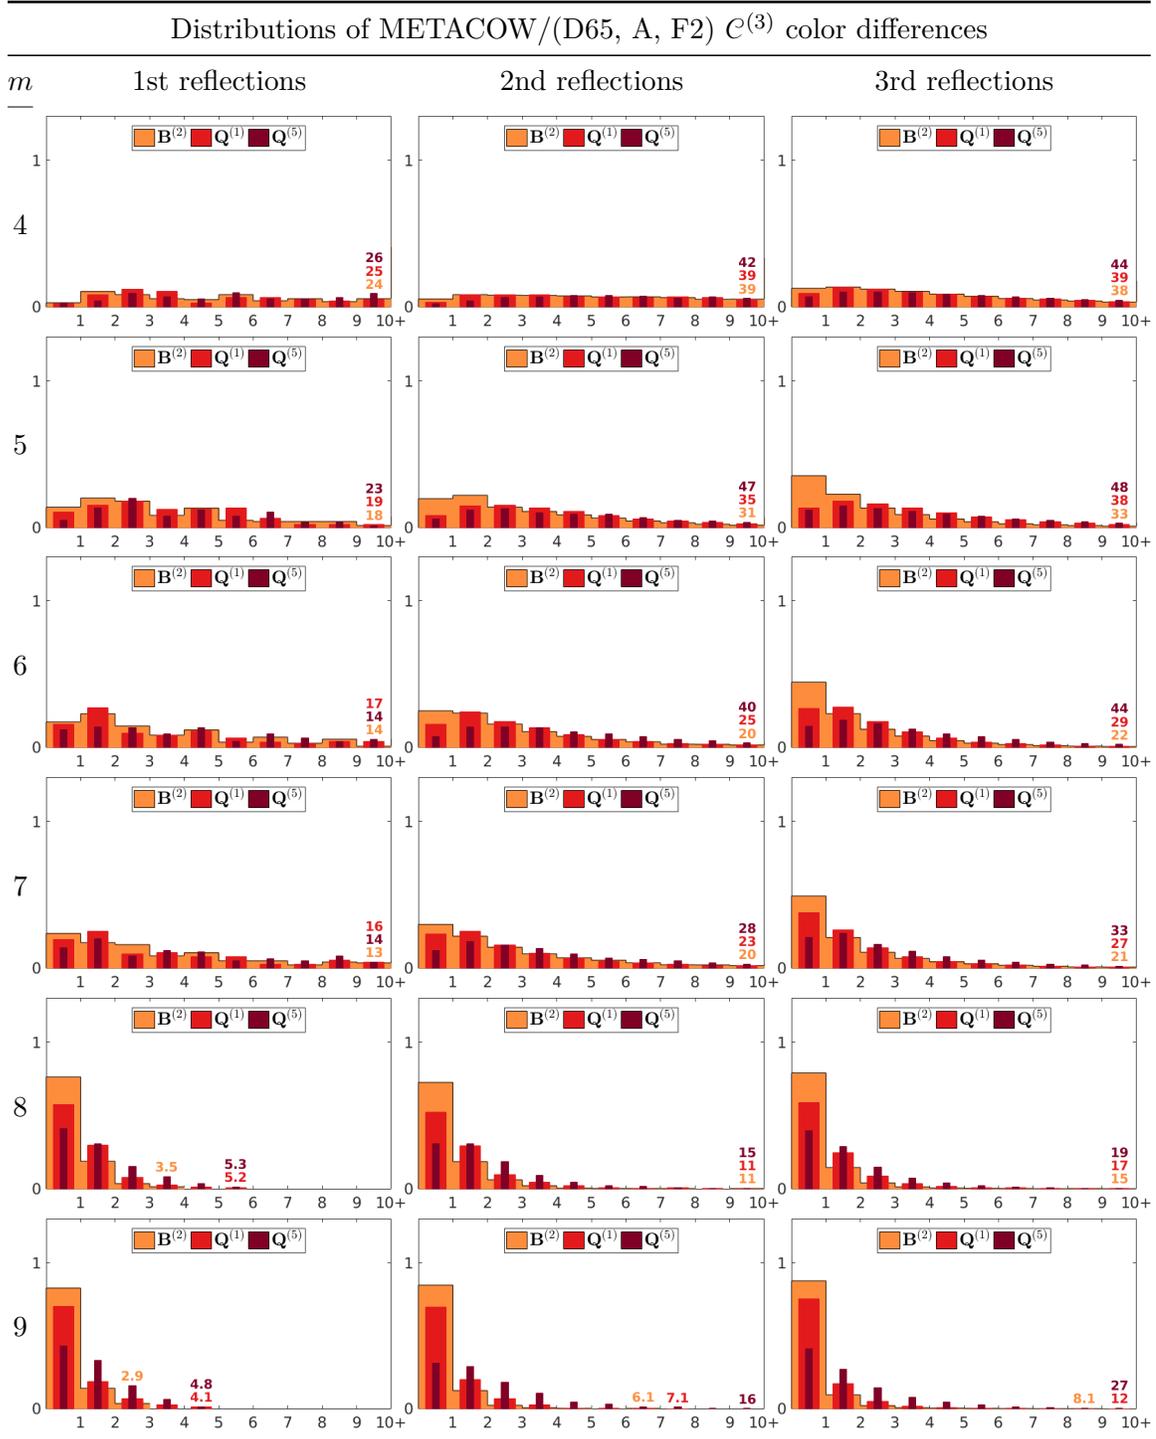


Figure A36: Distributions of ΔE_{00} color differences (as fractions of the totals) incurred when computing METACOW/(D65, A, F2) $C^{(3)}$ color signals by multiplying coefficients of $Q^{(1)}$ or $Q^{(5)}$, for dimensionalities $m \equiv 4, \dots, 9$ and first, second and third reflections. For comparison, differences incurred by projection onto $B^{(2)}$ are also shown. Maximum values, rounded to two significant figures, are shown in corresponding colors.

Bibliography

- [1] ABDUL-RAHMAN, A., AND CHEN, M. Spectral volume rendering based on the Kubelka-Munk theory. *Computer Graphics Forum* 24, 3 (2005), 413–422.
- [2] AGAHIAN, F., FUNT, B., AND AMIRSHAHI, S. H. Spectral compression: Weighted principal component analysis versus weighted least squares. In *Proceedings of the SPIE* (2014), B. E. Rogowitz, T. N. Pappas, and H. de Ridder, Eds., vol. 9014, Human Vision and Electronic Imaging XIX, pp. 90140Z–1–90140Z–6.
- [3] ANDERSON, M., MOTTA, R., CHANDRASEKAR, S., AND STOKES, M. Proposal for a standard default color space for the Internet – sRGB. In *Proceedings of the IS&T/SID 4th Color and Imaging Conference* (1996), Society for Imaging and Technology, pp. 238–245.
- [4] ARNOLD, S. E., FARUQ, S., SAVOLAINEN, V., MCOWAN, P. W., AND CHITTKA, L. FReD: The floral reflectance database – a Web portal for analyses of flower colour. *PLoS ONE* 5, 12:e14287 (2010), 1–9.
- [5] ASTM. Standard practice for computing the colors of objects by using the CIE system. Tech. Rep. E 308-08, American Society for Testing and Materials, 2009.
- [6] BARAKZEHI, M., AMIRSHAHI, S. H., PEYVANDI, S., AND AFJEH, M. G. Reconstruction of total radiance spectra of fluorescent samples by means of non-linear principal component analysis. *Journal of the Optical Society of America A* 30, 9 (2013), 1862–1870.
- [7] BARANOSKI, G. V., ROKNE, J. G., AND XU, G. Virtual spectrophotometric measurements for biologically and physically-based rendering. In *Proceedings of*

- the Eighth Pacific Conference on Computer Graphics and Applications* (2000), IEEE, pp. 398–399.
- [8] BENEŠ, J., HARVEY, T. A., VELÁZQUEZ-ARMENDÁRIZ, E., AND PRUM, R. O. Visualization of color as birds see it. In *SIGGRAPH Asia 2013 Posters* (New York, NY, USA, 2013), SA '13, ACM, pp. 21:1–21:1.
- [9] BERGNER, S., MOLLER, T., TORY, M., AND DREW, M. A practical approach to spectral volume rendering. *IEEE Transactions on Visualization and Computer Graphics* 11, 2 (2005), 207–216.
- [10] BERTHIER, S. *Iridescences, les couleurs physiques des insectes*. Springer-Verlag, Paris, France, 2003.
- [11] BEYLKIN, G. On the fast algorithm for multiplication of functions in the wavelet bases. In *Proceedings of the International Conference "Wavelets and Applications"* (1993), Y. Meyer and S. Roques, Eds., pp. 53–61.
- [12] BONNARDEL, V., AND MALONEY, L. T. Daylight, biochrome surfaces, and human chromatic response in the Fourier domain. *Journal of the Optical Society of America A* 17, 4 (2000), 677–686.
- [13] BORGES, C. F. Trichromatic approximation method for surface illumination. *Journal of the Optical Society of America A* 8, 8 (1991), 1319–1323.
- [14] BRILL, M. H. A non-PC look at principal components. *Color Research & Application* 28, 1 (2003), 69–71.
- [15] BROYDEN, C. G. The convergence of a class of double-rank minimization algorithms 1, parts I and II. *IMA Journal of Applied Mathematics* 6, 1 (1970), 76–90.
- [16] BRUCE, L. M., KOGER, C. H., AND LI, J. Dimensionality reduction of hyperspectral data using discrete wavelet transform feature extraction. *IEEE Transactions on geoscience and remote sensing* 40, 10 (2002), 2331–2338.
- [17] BUCHSBAUM, G., AND BLOCH, O. Color categories revealed by non-negative matrix factorization of Munsell color spectra. *Vision research* 42, 5 (2002), 559–563.

- [18] BUSSLER, M., ERTL, T., AND SADLO, F. Photoelasticity raycasting. *Computer Graphics Forum* 34, 3 (2015), 141–150.
- [19] CALLET, P. 3D reconstruction from 3D cultural heritage models. In *3D Research Challenges in Cultural Heritage: A Roadmap in Digital Heritage Preservation*, M. Ioannides and E. Quak, Eds., Lecture Notes in Computer Science 8355. Springer, Berlin, Heidelberg, 2014, pp. 135–142.
- [20] CAO, Q., WAN, X., LI, J., LIU, Q., LIANG, J., AND LI, C. Spectral data compression using weighted principal component analysis with consideration of human visual system and light sources. *Optical Review* 23, 5 (2016), 753–764.
- [21] CARDOSO, J.-F., AND SOULOUMIAC, A. Jacobi angles for simultaneous diagonalization. *Siam Journal on Matrix Analysis and Applications* 17, 1 (1996), 161–164.
- [22] CARMONA, P. L., AND LENZ, R. Performance evaluation of dimensionality reduction techniques for multispectral images. *International Journal of Imaging Systems and Technology* 17, 3 (2007), 202–217.
- [23] CARROLL, J. B. An analytical solution for approximating simple structure in factor analysis. *Psychometrika* 18, 1 (1953), 23–38.
- [24] CHABRIEL, G., KLEINSTEUBER, M., MOREAU, E., SHEN, H., TICHAVSKY, P., AND YEREDOR, A. Joint matrices decompositions and blind source separation: A survey of methods, identification, and applications. *IEEE Signal Processing Magazine* 31, 3 (2014), 34–43.
- [25] CHANE, C. S., MANSOURI, A., MARZANI, F. S., AND BOOCHS, F. Integration of 3D and multispectral data for cultural heritage applications: Survey and perspectives. *Image and Vision Computing* 31, 1 (2013), 91–102.
- [26] CHELLE, M., AND ANDRIEU, B. Radiative models for architectural modeling. *Agronomie* 19, 3–4 (1999), 225–240.
- [27] CHEN, Q., WANG, L. J., AND WESTLAND, S. The perceptual study of the tolerance of spectral images based on bootstrap analysis. *Advanced Materials Research* 301–303 (2011), 1151–1156.

- [28] CHEN, T. F., AND BARANOSKI, G. V. Effective compression and reconstruction of human skin hyperspectral reflectance databases. In *Proceedings of the 37th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)* (2015), IEEE, pp. 7027–7030.
- [29] CHEN, T. F., BARANOSKI, G. V., KIMMEL, B. W., AND MIRANDA, E. Hyperspectral modeling of skin appearance. *ACM Transactions on Graphics* 34, 3 (2015), 31:1–31:14.
- [30] CHERN, J.-R., AND WANG, C.-M. A novel progressive refinement algorithm for full spectral rendering. *Real-Time Imaging* 11, 2 (2005), 117–127.
- [31] CHIANG, M. J.-Y., KUTZ, P., AND BURLEY, B. Practical and controllable subsurface scattering for production path tracing. In *ACM SIGGRAPH 2016 Talks* (2016), SIGGRAPH '16, ACM, pp. 49:1–49:2.
- [32] CHIAO, C.-C., CRONIN, T. W., AND OSORIO, D. Color signals in natural scenes: characteristics of reflectance spectra and effects of natural illuminants. *Journal of the Optical Society of America A* 17, 2 (2000), 218–224.
- [33] CIE. Improvement to industrial colour-difference evaluation. Tech. Rep. Publication 142: 2001, Central Bureau of the Commission Internationale de l’Eclairage, 2001.
- [34] CIE. *Colorimetry, 3rd edition*, 2004. Publication 15: 2004.
- [35] CLAUSTRES, L., PAULIN, M., AND BOUCHER, Y. A wavelet-based framework for acquired radiometric quantity representation and accurate physical rendering. *The Visual Computer* 22, 4 (2006), 221–237.
- [36] COHEN, J. Dependency of the spectral reflectance curves of the Munsell color chips. *Psychonomic Science* 1, 1–12 (1964), 369–370.
- [37] COMON, P., AND JUTTEN, C. *Handbook of Blind Source Separation: Independent Component Analysis and Applications*. Academic Press, Burlington, Massachusetts, 2010.
- [38] CORMAN, T. H., LEISERSON, C. E., RIVET, R. L., AND STEIN, C. *Introduction to Algorithms*, 3rd ed. MIT Press, Cambridge, MA, 2009.

- [39] CRAWFORD, M. M., MA, L., AND KIM, W. Exploring nonlinear manifold learning for classification of hyperspectral data. In *Optical Remote Sensing: Advances in Signal Processing and Exploitation Techniques*, S. Prasad, L. M. Bruce, and J. Chanussot, Eds. Springer, Berlin, Heidelberg, 2011, pp. 207–234.
- [40] DANNEMILLER, J. L. Spectral reflectance of natural objects: How many basis functions are necessary? *Journal of the Optical Society of America A* 9, 4 (1992), 507–515.
- [41] DARLING, B. A. *A Framework for Realistic Modeling and Display of Object Surface Appearance*. PhD thesis, Rochester Institute of Technology, 2013.
- [42] DARLING, B. A., AND FERWERDA, J. A. Seeing virtual objects: simulating reflective surfaces on emissive displays. In *Proceedings of the IS&T/SID 20th Color and Imaging Conference (2012)*, Society for Imaging Science and Technology, pp. 135–141.
- [43] DARLING, B. A., FERWERDA, J. A., BERNS, R. S., AND CHEN, T. Real-time multispectral rendering with complex illumination. In *Proceedings of the IS&T/SID 19th Color and Imaging Conference (2011)*, Society for Imaging Science and Technology, pp. 345–351.
- [44] DEEB, R., MUSELET, D., HEBERT, M., AND TREMEAU, A. Interreflections in computer vision: A survey and an introduction to spectral infinite-bounce model. *Journal of Mathematical Imaging and Vision* (2017), 1–20.
- [45] DELAHUNT, P. B., AND BRAINARD, D. H. Color constancy under changes in reflected illumination. *Journal of Vision* 4, 9 (2004), 764–778.
- [46] DEVILLE, P. M., MERZOUK, S., CAZIER, D., AND PAUL, J. C. Spectral data modeling for a lighting application. *Computer Graphics Forum* 13, 3 (1994), C97–C106.
- [47] DISNEY, M., LEWIS, P., AND NORTH, P. Monte Carlo ray tracing in optical canopy reflectance modelling. *Remote Sensing Reviews* 18, 2–4 (2000), 163–196.
- [48] DREW, M. S., AND FINLAYSON, G. D. Spectral sharpening with positivity. *Journal of the Optical Society of America A* 17, 8 (2000), 1361–1370.

- [49] DREW, M. S., AND FINLAYSON, G. D. Multispectral processing without spectra. *Journal of the Optical Society of America A* 20, 7 (2003), 1181–1193.
- [50] DREW, M. S., AND FUNT, B. V. Natural metamers. *CVGIP: Image Understanding* 56, 2 (1992), 139–151.
- [51] EBERHART, R., AND KENNEDY, J. A new optimizer using particle swarm theory. In *Proceedings of the Sixth International Symposium on Micro Machine and Human Science* (1995), IEEE, pp. 39–43.
- [52] EVANS, G., AND MCCOOL, M. D. Stratified wavelength clusters for efficient spectral monte carlo rendering. In *Proceedings of Graphics Interface* (1999), vol. 99.
- [53] FAIRCHILD, M. D. *Color Appearance Models*, third ed. John Wiley & Sons, New York, NY, 2013.
- [54] FAIRCHILD, M. D., AND JOHNSON, G. M. METACOW: A public-domain, high-resolution, fully-digital, noise-free, metameric, extended-dynamic-range, spectral test target for imaging system analysis and simulation. In *Proceedings of the IS&T/SID 12th Color and Imaging Conference* (2004), Society for Imaging Science and Technology, pp. 239–245.
- [55] FAIRMAN, H. S., AND BRILL, M. H. The principal components of reflectances. *Color Research & Application* 29, 2 (2004), 104–110.
- [56] FAIRMAN, H. S., BRILL, M. H., AND HEMMENDINGER, H. How the CIE 1931 color-matching functions were derived from Wright-Guild data. *Color Research & Application* 22, 1 (1997), 11–23.
- [57] FARAJIKHAH, S., MADANCHI, F., AND AMIRSHAHI, S. H. Nonlinear principal component analysis for compression of spectral data. In *Proceedings of Conference on Data Mining and Data Warehouses* (2011).
- [58] FARRELL, J. E., XIAO, F., CATRYSSSE, P. B., AND WANDELL, B. A. A simulation tool for evaluating digital camera image quality. In *Image Quality and System Performance* (2003), vol. SPIE 5294, pp. 124–132.
- [59] FERGUSON, G. A. The concept of parsimony in factor analysis. *Psychometrika* 19, 4 (1954), 281–290.

- [60] FLETCHER, R. A new approach to variable metric algorithms. *The Computer Journal* 13, 3 (1970), 317–322.
- [61] FLETCHER, R. *Practical Methods of Optimization*, second ed. John Wiley & Sons, New York, NY, 1987.
- [62] FU, Y., TIPPETS, C. A., DONEV, E. U., AND LOPEZ, R. Structural colors: from natural to artificial systems. *Wiley Interdisciplinary Reviews: Nanomedicine and Nanobiotechnology* 8, 5 (2016), 758–775.
- [63] FUNT, B., KULPINSKI, D., AND CARDEI, V. Non-linear embeddings and the underlying dimensionality of reflectance spectra and chromaticity histograms. In *Proceedings of the IS&T/SID 9th Color and Imaging Conference* (2001), Society for Imaging Science and Technology, pp. 126–129.
- [64] GARCÍA-BELTRÁN, A., NIEVES, J. L., HERNÁNDEZ-ANDRÉS, J., AND ROMERO, J. Linear bases for spectral reflectance functions of acrylic paints. *Color Research & Application* 23, 1 (1998), 39–45.
- [65] GEISLER-MORODER, D., AND DÜR, A. Color-rendering indices in global illumination methods. *Journal of Electronic Imaging* 18, 4 (2009), 043015–1–043015–12.
- [66] GEIST, R., HEIM, O., AND JUNKINS, S. Color representation in virtual environments. *Color Research & Application* 21, 2 (1996), 121–128.
- [67] GLASSNER, A. S. *Principles of digital image synthesis*. Morgan Kaufmann, San Francisco, CA, 1995.
- [68] GOLDFARB, D. A family of variable-metric methods derived by variational means. *Mathematics of Computation* 24, 109 (1970), 23–26.
- [69] GOLUB, G. H., AND LOAN, C. F. V. *Matrix Computations*, fourth ed. Johns Hopkins University Press, Baltimore, 2013.
- [70] GOVAERTS, Y. M., AND VERSTRAETE, M. M. Raytran: A Monte Carlo ray-tracing model to compute light scattering in three-dimensional heterogeneous media. *IEEE Transactions on geoscience and remote sensing* 36, 2 (1998), 493–505.

- [71] HALL, R. *Illumination and color in computer generated imagery*. Springer Verlag, New York, NY, 1989.
- [72] HALL, R. Comparing spectral color computation methods. *IEEE Computer Graphics and Applications* 19, 4 (1999), 36–45.
- [73] HAPPA, J., BASHFORD-ROGERS, T., WILKIE, A., ARTUSI, A., DEBATTISTA, K., AND CHALMERS, A. Cultural heritage predictive rendering. *Computer Graphics Forum* 31, 6 (2012), 1823–1836.
- [74] HAUTA-KASARILL, M., WANG, W., TOYOOKA, S., PARKKINEN, J., AND LENZ, R. Unsupervised filtering of Munsell spectra. In *Asian Conference on Computer Vision* (1998), Springer, pp. 248–255.
- [75] HEALEY, G., AND BENITES, L. Linear models for spectral reflectance functions over the mid-wave and long-wave infrared. *Journal of the Optical Society of America A* 15, 8 (1998), 2216–2227.
- [76] HORST, P. *Factor analysis of data matrices*. Holt, Rinehart and Winston, New York, NY, 1965.
- [77] HUANG, R., QIAO, W., YANG, J., WANG, H., XUE, B., AND TAO, J. Multispectral image compression algorithm based on spectral clustering and wavelet transform. In *LIDAR Imaging Detection and Target Recognition 2017* (2017), vol. 10605, International Society for Optics and Photonics, p. 106051X.
- [78] IEHL, J. C., AND PÉROCHE, B. An adaptive spectral rendering with a perceptual control. *Computer Graphics Forum* 19, 3 (2000), 291–299.
- [79] IEHL, J. C., AND PÉROCHE, B. Towards perceptual control of physically based spectral rendering. *Computers & Graphics* 27, 5 (2003), 747–762.
- [80] IMAI, F. H., ROSEN, M. R., AND BERNS, R. S. Comparative study of metrics for spectral match quality. In *CGIV 2002: The First European Conference on Colour Graphics, Imaging, and Vision* (2002), Society for Imaging Science and Technology, pp. 492–496.
- [81] ISO/CIE. *Colorimetry – Part 2: CIE Standard Illuminants for Colorimetry*, 2007. Joint Standard ISO 11664-2:2007(E) / CIE S 014-2/E:2006.

- [82] JAASKELAINEN, T., PARKKINEN, J., AND TOYOOKA, S. Vector-subspace model for color representation. *Journal of the Optical Society of America A* 7, 4 (1990), 725–730.
- [83] JACQUES, S. L. Monte Carlo modeling of light transport in tissue (steady state and time of flight). In *Optical-Thermal response of laser-irradiated tissue*. Springer, 2010, pp. 109–144.
- [84] JENNRICH, R. I. A simple general procedure for orthogonal rotation. *Psychometrika* 66, 2 (2001), 289–306.
- [85] JENNRICH, R. I., AND SAMPSON, P. Rotation for simple loadings. *Psychometrika* 31, 3 (1966), 313–323.
- [86] JIANG, J., LIU, D., GU, J., AND SÜSTRUNK, S. What is the space of spectral sensitivity functions for digital color cameras? In *2013 IEEE Workshop on Applications of Computer Vision (WACV)* (2013), pp. 168–179.
- [87] JOHNSON, G. M., AND FAIRCHILD, M. D. Full-spectral color calculations in realistic image synthesis. *IEEE Computer Graphics and Applications* 19, 4 (1999), 47–53.
- [88] JOLLIFFE, I. T. *Principal component analysis*, 2nd ed. Springer-Verlag, New York, NY, 2002.
- [89] KAARNA, A. *Multispectral image compression using the wavelet transform*. PhD thesis, Lappeenranta University of Technology, 2000.
- [90] KAEWPIJIT, S., LE MOIGNE, J., AND EL-GHAZAWI, T. Automatic reduction of hyperspectral imagery using wavelet spectral analysis. *IEEE transactions on Geoscience and Remote Sensing* 41, 4 (2003), 863–871.
- [91] KAISER, H. F. The varimax criterion for analytic rotation in factor analysis. *Psychometrika* 23, 3 (1958), 187–200.
- [92] KARAMI, A., YAZDI, M., AND MERCIER, G. Compression of hyperspectral images using discrete wavelet transform and Tucker decomposition. *IEEE journal of selected topics in applied earth observations and remote sensing* 5, 2 (2012), 444–450.

- [93] KIM, C.-Y., SEO, Y.-S., AND KWEON, I.-S. Color linear model. In *Proceedings of the 9th International Conference on Image Analysis and Processing* (1997), Springer-Verlag, pp. 54–61.
- [94] KIM, M. H., HARVEY, T. A., KITTLE, D. S., RUSHMEIER, H., DORSEY, J., PRUM, R. O., AND BRADY, D. J. 3D imaging spectroscopy for measuring hyperspectral patterns on solid objects. *ACM Transactions on Graphics* 31, 4 (2012), 38:1–38:11.
- [95] KIM, M. H., RUSHMEIER, H., FFRENCH, J., PASSERI, I., AND TIDMARSH, D. Hyper3D: 3D graphics software for examining cultural artifacts. *Journal on Computing and Cultural Heritage* 7, 3 (2014), 14:1–14:19.
- [96] KINOSHITA, S. *Structural colors in the realm of nature*. World Scientific, Singapore, 2008.
- [97] KIRKPATRICK, S., GELATT JR., C. D., AND VECCHI, M. P. Optimization by simulated annealing. *Science* 220, 4598 (1983), 671–680.
- [98] KOHONEN, O., PARKKINEN, J., AND JÄÄSKELÄINEN, T. Databases for spectral color science. *Color Research and Application* 31, 5 (2006), 381–390.
- [99] KUTZ, P., HABEL, R., LI, Y. K., AND NOVÁK, J. Spectral and decomposition tracking for rendering heterogeneous volumes. *ACM Transactions on Graphics* 36, 4 (2017), 111:1–111:16.
- [100] LAAMANEN, H., JAEAESKELAEINEN, T., AND PARKKINEN, J. P. S. Comparison of PCA and ICA in color recognition. In *Proceedings of Intelligent Robots and Computer Vision XIX: Algorithms, Techniques, and Active Vision* (2000), vol. SPIE 4197, pp. 367–377.
- [101] LAAMANEN, H., JETSU, T., JAASKELAINEN, T., AND PARKKINEN, J. P. S. Weighted compression of spectral color information. *Journal of the Optical Society of America A* 25, 6 (2008), 1383–1388.
- [102] LAM, M. W., AND BARANOSKI, G. V. A predictive light transport model for the human iris. *Computer Graphics Forum* 25, 3 (2006), 359–368.

- [103] LAWRENCE, N. D. A unifying probabilistic perspective for spectral dimensionality reduction: Insights and new models. *Journal of Machine Learning Research* 13, May (2012), 1609–1638.
- [104] LEHTONEN, J., PARKKINEN, J., AND JAASKELAINEN, T. Optimal sampling of color spectra. *Journal of the Optical Society of America A* 23, 12 (2006), 2983–2988.
- [105] LEHTONEN, J., PARKKINEN, J., JAASKELAINEN, T., AND KAMSHILIN, A. Principal component and sampling analysis of color spectra. *Optical Review* 16, 2 (2009), 81–90.
- [106] LENZ, R., AND BUI, T. H. Statistical properties of color-signal spaces. *Journal of the Optical Society of America A* 22, 5 (2005), 820–827.
- [107] LENZ, R., HILTUNEN, J., PARKKINEN, J., JAASKELAINEN, T., AND ÖSTERBERG, M. Unsupervised filtering of color spectra. *Journal of the Optical Society of America A* 13, 7 (1996), 1315–1324.
- [108] LEVENBERG, K. A method for the solution of certain non-linear problems in least squares. *Quarterly of Applied Mathematics* 2, 2 (1944), 164–168.
- [109] LEVERENZ, H. W. *An introduction to luminescence of solids*. John Wiley & Sons, New York, NY, 1950.
- [110] LONG, J., ESTEY, A., BARTLE, D., OLSEN, S., AND GOOCH, A. A. Catalyst: seeing through the eyes of a cat. In *Proceedings of the Fifth International Conference on the Foundations of Digital Games* (2010), ACM, pp. 116–123.
- [111] LONG, J., AND GOOCH, A. A. Bee prepared: Simulating bee vision in an educational game. In *2011 16th International Conference on Computer Games (CGAMES)* (2011), pp. 262–269.
- [112] MAGNUS, J. R., AND NEUDECKER, H. *Matrix differential calculus with applications in statistics and econometrics*, third ed. John Wiley & Sons, New York, NY, 2007.
- [113] MAHY, M., EYCKEN, L. V., AND OOSTERLINCK, A. Evaluation of uniform color spaces developed after the adoption of CIELAB and CIELUV. *Color Research and Application* 19, 2 (1994), 105–1121.

- [114] MALLAT, S. *A Wavelet Tour of Signal Processing: The Sparse Way*, third ed. Academic Press, Burlington, MA, 2008.
- [115] MALONEY, L. T. Evaluation of linear models of surface spectral reflectance with small numbers of parameters. *Journal of the Optical Society of America A* 3, 10 (1986), 1673–1683.
- [116] MALONEY, L. T. Physics-based approaches to modeling surface color perception. In *Color vision: From genes to perception*, K. R. Gegenfurtner and L. T. Sharpe, Eds. Cambridge University Press, 1999, pp. 387–416.
- [117] MANSOURI, A., SLIWA, T., HARDEBERG, J. Y., AND VOISIN, Y. Representation and estimation of spectral reflectances using projection on PCA and wavelet bases. *Color Research & Application* 33, 6 (2008), 485–493.
- [118] MARIMONT, D. H., AND WANDELL, B. A. Linear models of surface and illuminant spectra. *Journal of the Optical Society of America A* 9, 11 (1992), 1905–1913.
- [119] MARKAS, T., AND REIF, J. Multispectral image compression algorithms. In *Proceedings of the Data Compression Conference (DCC'93)* (1993), IEEE, pp. 391–400.
- [120] MARQUARDT, D. W. An algorithm for least-squares estimation of nonlinear parameters. *Journal of the Society for Industrial and Applied Mathematics* 11, 2 (1963), 431–441.
- [121] MCCAMY, C. S., MARCUS, H., AND DAVIDSON, J. A color-rendition chart. *Journal of Applied Photographic Engineering* 2, 3 (1976), 95–99.
- [122] MESETH, J. *Towards Predictive Rendering in Virtual Reality*. PhD thesis, University of Bonn, 2006.
- [123] MEYER, G. W. Wavelength selection for synthetic image generation. *Computer Vision, Graphics, and Image Processing* 41, 1 (1988), 57–79.
- [124] MOON, P., AND SPENCER, D. E. Polynomial representation of reflectance curves. *Journal of the Optical Society of America A* 35, 9 (1945), 597–600.

- [125] MOROVIČ, J. *Color Gamut Mapping*. John Wiley & Sons, New York, NY, 2008.
- [126] MULAİK, S. A. *Foundations of Factor Analysis*. CRC press, Boca Raton, FL, 2009.
- [127] MULLER, T., CALLET, P., DA GRAÇA, F., PALJIC, A., PORRAL, P., AND HOARAU, R. Predictive rendering of composite materials: a multi-scale approach. In *Measuring, Modeling, and Reproducing Material Appearance 2015* (2015), vol. SPIE 9398, pp. 9398–2–9398–14.
- [128] MUNSELL COLOR. *Munsell Book of Color – Matte Finish Collection*. Baltimore, MD, 1976.
- [129] NASCIMENTO, S. M., FOSTER, D. H., AND AMANO, K. Psychophysical estimates of the number of spectral-reflectance basis functions needed to reproduce natural scenes. *Journal of the Optical Society of America A* 22, 6 (2005), 1017–1022.
- [130] NASSAU, K. *The Physics and Chemistry of Color: The Fifteen Causes of Color*, second ed. John Wiley & Sons, New York, NY, 2001.
- [131] NĚMCOVÁ, Š., HAVRAN, V., ČÁP, J., HOŠEK, J., AND BITTNER, J. Compensating for color artifacts in the design of technical kaleidoscopes. *Optical Engineering* 55, 11 (2016), 113104–1–113104–7.
- [132] NEUHAUS, J. O., AND WRIGLEY, C. The quartimax method. *British Journal of Statistical Psychology* 7, 2 (1954), 81–91.
- [133] NOCEDAL, J., AND WRIGHT, S. J. *Numerical optimization*, second ed. Springer, New York, NY, 2006.
- [134] NORTH, P. R. Three-dimensional forest light interaction model using a Monte Carlo method. *IEEE transactions on Geoscience and Remote Sensing* 34, 4 (1996), 946–956.
- [135] OKADA, N., ZHU, D., CAI, D., COLE, J. B., KAMBE, M., AND KINOSHITA, S. Rendering *Morpho* butterflies based on high accuracy nano-optical simulation. *Journal of Optics* 42, 1 (2013), 25–36.

- [136] OWENS, H. C., DOHERTY, L., AND WESTLAND, S. Representation of natural reflectance spectra by auto-associative neural network. In *Colour Image Science 2000 (Derby) Proceedings* (2000), pp. 304–311.
- [137] OXTOBY, E. K., AND FOSTER, D. H. Perceptual limits on low-dimensional models of Munsell reflectance spectra. *Perception* 34, 8 (2005), 961–966.
- [138] PAN, V. Y. How bad are Vandermonde matrices? *SIAM Journal on Matrix Analysis and Applications* 37, 2 (2016), 676–694.
- [139] PARKKINEN, J. P., HALLIKAINEN, J., AND JAASKELAINEN, T. Characteristic spectra of Munsell colors. *Journal of the Optical Society of America A* 6, 2 (1989), 318–322.
- [140] PEERCY, M. S. Linear color representations for full speed spectral rendering. In *Proceedings of the 20th Annual Conference on Computer Graphics and Interactive Techniques* (1993), SIGGRAPH '93, pp. 191–198.
- [141] PEINECKE, N. The challenges of closed-loop sensor simulation. *IEEE Aerospace and Electronic Systems Magazine* 29, 6 (2014), 37–43.
- [142] PENNA, B., TILLO, T., MAGLI, E., AND OLMO, G. Transform coding techniques for lossy hyperspectral data compression. *IEEE Transactions on Geoscience and Remote Sensing* 45, 5 (2007), 1408–1421.
- [143] PREECE, S. J., AND CLARIDGE, E. Monte Carlo modelling of the spectral reflectance of the human eye. *Physics in Medicine and Biology* 47, 16 (2002), 2863–2877.
- [144] RAAB, M., SEIBERT, D., AND KELLER, A. Unbiased global illumination with participating media. In *Monte Carlo and Quasi-Monte Carlo Methods*. Springer, 2008, pp. 591–605.
- [145] RADZISZEWSKI, M., BORYCZKO, K., AND ALDA, W. An improved technique for full spectral rendering. *Journal of WSCG* 17, 1–3 (2009), 9–16.
- [146] RAMANATH, R., KUEHNI, R. G., SNYDER, W. E., AND HINKS, D. Spectral spaces and color spaces. *Color Research & Application* 29, 1 (2004), 29–37.

- [147] RASO, M. G., AND FOURNIER, A. A piecewise polynomial approach to shading using spectral distributions. In *Proceedings of Graphics Interface '91* (1991), pp. 40–46.
- [148] REINHARD, E., HEIDRICH, W., DEBEVEC, P., PATTANAİK, S., WARD, G., AND MYZKOWSKI, K. *High Dynamic Range Imaging: Acquisition, Display, and Image-Based Lighting*. Morgan Kaufmann, Burlington, MA, 2010.
- [149] RODARMEL, C., AND SHAN, J. Principal component analysis for hyperspectral image classification. *Surveying and Land Information Science* 62, 2 (2002), 115–122.
- [150] ROMERO, J., GARCIA-BELTRÁN, A., AND HERNÁNDEZ-ANDRÉS, J. Linear bases for representation of natural and artificial illuminants. *Journal of the Optical Society of America A* 14, 5 (1997), 1007–1014.
- [151] ROMERO, J., VALERO, E., HERNÁNDEZ-ANDRÉS, J., AND NIEVES, J. L. Color-signal filtering in the Fourier-frequency domain. *Journal of the Optical Society of America A* 20, 9 (2003), 1714–1724.
- [152] ROUGERON, G., AND PÉROCHE, B. An adaptive representation of spectral data for reflectance computations. In *Rendering Techniques '97: Proceedings of the 8th Eurographics Rendering Workshop* (1997), J. Dorsey and P. Slusallek, Eds., Springer-Verlag, pp. 127–138.
- [153] SAUNDERS, D. R. An analytic method for rotation to orthogonal simple structure. Tech. Rep. Research Bulletin 53-10, Educational Testing Service, 1953.
- [154] SHANNO, D. F. Conditioning of quasi-Newton methods for function minimization. *Mathematics of Computation* 24, 111 (1970), 647–656.
- [155] SHERIN, R. J. A matrix formulation of Kaiser's varimax criterion. *Psychometrika* 31, 4 (1966), 535–538.
- [156] STOKES, M., FAIRCHILD, M. D., AND BERNS, R. S. Precision requirements for digital color reproduction. *ACM Transactions on Graphics* 11, 4 (1992), 406–422.

- [157] SUBRAMANIAN, S., GAT, N., RATCLIFF, A., AND EISMANN, M. Real-time hyperspectral data compression using principal components transformation. In *Proceedings of the AVIRIS Earth Science & Applications Workshop* (2000).
- [158] SUN, Y., FRACCHIA, F. D., DREW, M. S., AND CALVERT, T. W. A spectrally based framework for realistic image synthesis. *The Visual Computer* 17, 7 (2001), 429–444.
- [159] TALBI, E.-G. *Metaheuristics: From design to implementation*. John Wiley & Sons, Hoboken, New Jersey, 2009.
- [160] TANDIANUS, B., JOHAN, H., SEAH, H. S., AND LIN, F. Spectral caustic rendering of a homogeneous caustic object based on wavelength clustering and eye sensitivity. *The Visual Computer* 31, 12 (2015), 1601–1614.
- [161] TICHAVSKÝ, P., YEREDOR, A., AND NIELSEN, J. A fast approximate joint diagonalization algorithm using a criterion with a block diagonal weight matrix. In *2008 International Conference on Acoustics, Speech and Signal Processing* (2008), IEEE, pp. 3321–3324.
- [162] TILLEY, R. J. D. *Colour and optical properties of materials: An exploration of the relationship between light, the optical properties of materials and colour*, second ed. J. Wiley & Sons, Chichester, West Sussex, U.K., 2011.
- [163] TRUSSELL, H. Applications of set theoretic methods to color systems. *Color Research & Application* 16, 1 (1991), 31–41.
- [164] TZENG, D.-Y., AND BERNS, R. S. A review of principal component analysis and its applications to color technology. *Color Research & Application* 30, 2 (2005), 84–98.
- [165] USUI, S., NAKAUCHI, S., AND NAKANO, M. Reconstruction of Munsell color space by a five-layer neural network. *Journal of the Optical Society of America A* 9, 4 (1992), 516–520.
- [166] VEACH, E. *Robust Monte Carlo methods for light transport simulation*. PhD thesis, Stanford University, 1998.

- [167] VIGGIANO, J. A. S. Metrics for evaluating spectral matches: A quantitative comparison. In *CGIV 2004: The Second European Conference on Colour Graphics, Imaging and Vision* (2004), Society for Imaging Science and Technology, pp. 286–291.
- [168] VRHEL, M. J., GERSHON, R., AND IWAN, L. S. Measurement and analysis of object reflectance spectra. *Color Research & Application* 19, 1 (1994), 4–9.
- [169] WACHTLER, T., LEE, T.-W., AND SEJNOWSKI, T. J. Chromatic structure of natural scenes. *Journal of the Optical Society of America A* 18, 1 (2001), 65–77.
- [170] WADDLE, C. Real-time spectral rendering. In *Proceedings of the IS&T/SID 21st Color and Imaging Conference* (2013), Society for Imaging and Technology, pp. 240–246.
- [171] WANDELL, B. A. The synthesis and analysis of color images. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 9, 1 (1987), 2–13.
- [172] WARD, G., AND EYDELBERG-VILESHIN, E. Picture perfect RGB rendering using spectral prefiltering and sharp color primaries. In *Thirteenth Eurographics Workshop on Rendering* (2002), P. Debevec and S. Gibson, Eds., pp. 117–124.
- [173] WEISSKOPF, V. F. How light interacts with matter. *Scientific American* 219, 3 (1968), 60–71.
- [174] WILKIE, A., NAWAZ, S., DROSKE, M., WEIDLICH, A., AND HANIKA, J. Hero wavelength spectral sampling. *Computer Graphics Forum* 33, 4 (2014), 123–131.
- [175] WILKIE, A., TOBLER, R. F., AND PURGATHOFER, W. Raytracing of dispersion effects in transparent materials. In *WSCG 2000 Conference Proceedings* (2000), Václav Skala-UNION Agency.
- [176] WILKIE, A., WEIDLICH, A., MAGNOR, M., AND CHALMERS, A. Predictive rendering. In *ACM SIGGRAPH ASIA 2009 Courses* (2009), SIGGRAPH ASIA '09, ACM, pp. 12:1–12:428.

- [177] WYSZECKI, G., AND STILES, W. *Color Science: Concepts and Methods, Quantitative Data and Formulae*, second ed. John Wiley & Sons, New York, NY, 1982.
- [178] XIAO, Z., AND BOURENNANE, S. Constrained nonnegative matrix factorization and hyperspectral image dimensionality reduction. *Remote Sensing Letters* 5, 1 (2014), 46–54.
- [179] XIONG, W., AND FUNT, B. Independent component analysis and nonnegative linear model analysis of illuminant and reflectance spectra. In *Proceedings of the 10th Congress of the International Colour Association (AIC '05)* (2005), pp. 503–506.
- [180] XU, H., AND SUN, Y. Representing scattering functions with spherical harmonics of spectral fourier components. In *Computational Imaging* (2004), pp. 22–33.
- [181] ZEGHERS, E., CARRÉ, S., AND BOUATOUCH, K. Error-bound wavelength selection for spectral rendering. *The Visual Computer* 13, 9 (1998), 424–434.
- [182] ZHANG, X., AND WANDELL, B. A. A spatial extension of CIELAB for digital color-image reproduction. *Journal of the Society for Information Display* 5, 1 (1997), 61–63.
- [183] ZHU, C., AND LIU, Q. Review of Monte Carlo modeling of light transport in tissues. *Journal of Biomedical Optics* 18, 5 (2013), 050902–050902.