

Orientation Invariant Light Source Parameters

R. D. Stock and M. W. Siegel

Keywords: light source characterization, rotational invariants,
data compression, interpolation

Carnegie Mellon University
Measurement and Control Laboratory
5000 Forbes Avenue
Pittsburgh, PA 15213
phone: (412) 268-8802
fax: (412) 621-5477
e-mail: rstock+@sensor.ri.cmu.edu
e-mail: mws+@isl1.ri.cmu.edu

1 Abstract

At present there is no common language for comparing two different light sources. Trying to quantify (or even to explain qualitatively) the difference in performance terms, for example, of an axial and a transverse filament, or a gas discharge tube and a light emitting diode, or a gas lamp mantle and a laser, is like comparing apples and oranges. We have developed a concept for a common representation of these diverse kinds of light sources, without reference to any coordinate system or light generation process. Based on an expansion in spherical harmonics of an intensity dataset over a sphere, we characterize the angular light distribution of a source by providing a prescription for calculating the orientation independent parameters d_l that represent the proportion of *rms* power distributed in each angular mode.

2 Data Collection and Analysis

This luminaire model utilizes a goniometer to acquire intensity data equally spaced in co-latitude θ and longitude ϕ over a sphere of radius r_0 surrounding the light source. Light emanating from the source in a particular (radial) direction is collected by a lens, focused on a stop, and then falls on the surface of a photodiode. Figure 1 shows a simple schematic of the operation of a typical goniometer and detector apparatus. The purpose of the lens and stop is to collimate the light, eliminating extraneous signals and illuminating the photodiode over its entire surface from a mainly forward direction.¹ The photodiode integrates the incoming light, resulting in a single intensity value for each point on the sphere. Because of its simplicity, this dataset is useful for calculating rotationally invariant quantities that characterize or describe the light source. Such measurements have been used for raytracing purposes¹, but are only appropriate for that task in the far-zone.

The measured illumination distribution function² $I[\theta,\phi]$ is fit to a series of spherical harmonics

¹No attempt is made to account for the angular response of the photodiode, but the optics keep the angle of incidence within a narrow cone about the photodiode's surface normal.

²Throughout this paper, experimental parameters will be in square brackets and free variables will be in parentheses.

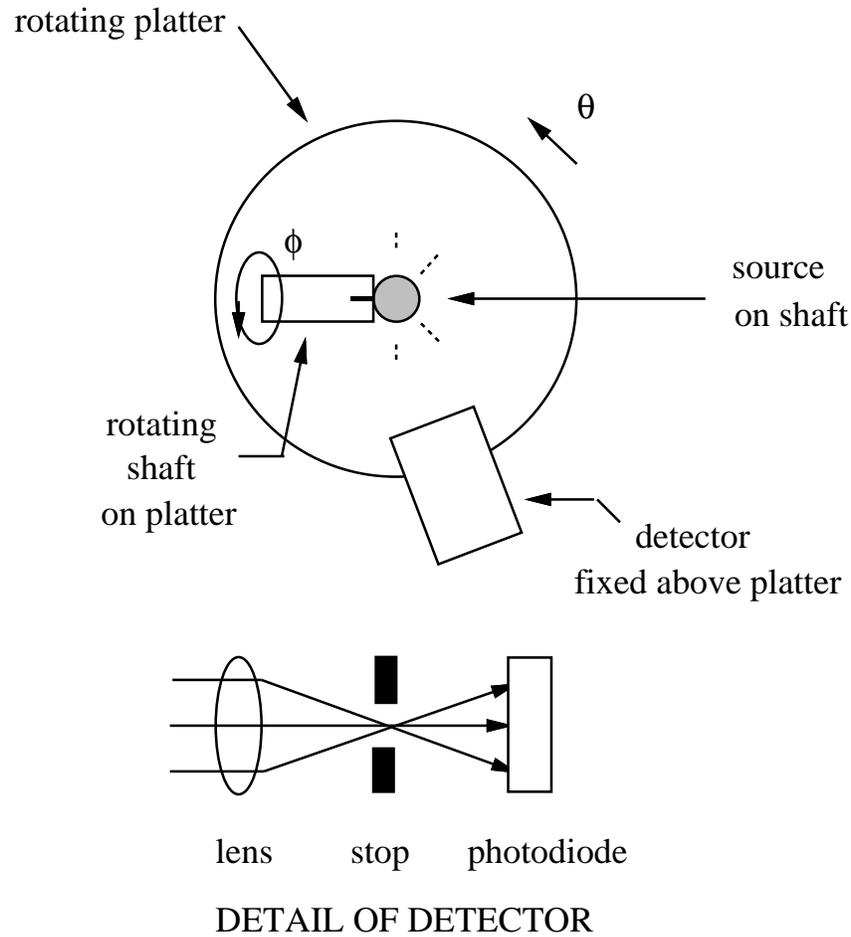


Figure 1: goniometer and detector apparatus

times a set of complex coefficients³ $\{\alpha_{lm}\}$:

$$I[\theta, \phi] = \sum_{lm} \alpha_{lm} Y_{lm}[\theta, \phi]$$

A comprehensive discussion of spherical harmonics is found in Varshalovich².

The coefficients are found directly by a discrete Fourier "dot product" integral:

$$\alpha_{lm} = \sum_i Y_{lm}^*[\theta, \phi]_i I[\theta, \phi]_i \sin \theta_i \Delta \theta \Delta \phi$$

where Y_{lm}^* is the complex conjugate of Y_{lm} .

Alternatively, the coefficients could be found by a least-squares fit. This has the advantage of built-in statistical tools³ that allow for the automatic computation of the uncertainties of the

³With real data, the appropriate range for m is 0 to l .

coefficients α_{lm} , although the computation is a little bit slower than the direct evaluation of the dot product. Issues concerning the data fitting process are discussed more fully in appendix II.

The coefficients α_{lm} themselves, however, cannot be used directly to quantitatively compare the performance qualities of different light sources; the same light source will produce a radically different set of coefficients if it is rotated with respect to the measuring apparatus. The problem then becomes one of finding rotational invariants among the coefficients to calculate numbers that capture the essence of the angular distribution of illumination produced by the source in question.

3 Orientation Independent Source Characterization

Because the Y_{lm} form an orthogonal basis in θ and ϕ , the α_{lm} can be thought of as the components of a vector whose length remains constant under rotations. Therefore,

$$\sum_{lm} |\alpha_{lm}|^2 = \text{constant} = N$$

That just says that the total mean-square intensity does not depend on the orientation of the light source about the center of the measurement sphere. But it is possible to go even further.

Consider now an arbitrary rotation of coordinates. The Y_{lm} transform as⁴

$$Y_{lm}(\theta', \phi') = \sum_n D_{nm}^l Y_{ln}(\theta, \phi)$$

where the \mathbf{D} are the Wigner \mathbf{D} -matrices used in the quantum mechanics of angular momentum⁴, and depend only on the Euler angles of the rotation. They possess two important properties. The first is that there is a different \mathbf{D} for each l -value. Thus a rotation can transform Y_{20} into some combination of Y_{21} and Y_{22} , but it can never be transformed into Y_{30} : rotations do not mix l -values.⁴ The second important fact is that the \mathbf{D} -matrices all have determinant +1 and are therefore length preserving in each l subspace. It follows then that for **each** l , the length of the projection of the "coefficient vector" α into the subspace l is rotationally invariant, or

$$\left\{ \sum_m |\alpha_{lm}|^2 / N \right\}^{1/2} = \text{constant} = d_l$$

⁴This is related to the fact that the Y_{lm} are the eigenfunctions of the angular momentum operator, where l depends on the total angular momentum and m is a component; rotations can change angular momentum components (m values), but cannot change the total angular momentum (l values).

A more detailed proof is given in appendix I.

The factor of N is included to normalize the d -numbers so that the squares of the d_l sum to unity, allowing them to be compared on an equal footing. The d_l , then, are rotationally independent quantities that can be used to characterize and compare the angular distributions of different light sources: d_0 represents the "monopole" strength (I_{avg}/I_{rms}), d_1 the "dipole" strength, and so on. Appendix III shows that physically the d_l represent the ratio of *rms* power distributed in the angular mode l to the total *rms* power.⁵ In the far-zone the normalized d_l become constant with radius as well as orientation, because the power (and hence all of the α_{lm}) is falling as $1/r^2$.

N and α_{00} have the following physical meanings. The mean-square intensity is defined as

$$\begin{aligned} I_{rms}^2 &= \int_{4\pi} d\Omega I(\Omega)^2 / 4\pi \\ &= \int_{4\pi} d\Omega \sum_{lm} \alpha_{lm} Y_{lm}(\Omega) \sum_{l'm'} \alpha_{l'm'} Y_{l'm'}(\Omega) / 4\pi \\ &= \sum_{ll'mm'} \alpha_{lm} \alpha_{l'm'} \int_{4\pi} d\Omega Y_{lm}(\Omega) Y_{l'm'}(\Omega) / 4\pi \\ &= \sum_{ll'mm'} \alpha_{lm} \alpha_{l'm'} \delta_{ll'} \delta_{mm'} / 4\pi \\ &= \sum_{lm} |\alpha_{lm}|^2 / 4\pi = N / 4\pi \end{aligned}$$

Thus

$$N = 4\pi I_{ms}^2$$

Furthermore, the average intensity is

$$\begin{aligned} I_{avg} &= \int_{4\pi} d\Omega I(\Omega) / 4\pi \\ &= \int_{4\pi} d\Omega I(\Omega) Y_{00}^2 = \alpha_{00} Y_{00} \end{aligned}$$

so that

$$\alpha_{00} = \sqrt{4\pi} I_{avg}$$

⁵In the quantum mechanical analogy, d_l^2 is the probability that the system has total orbital angular momentum l ; this probability should of course be independent of coordinate system rotations.

Another characteristic parameter can be calculated from the coefficients to represent the anisotropy, or standard deviation σ of the illumination distribution. Because $\alpha_{00}(4\pi)^{-1/2}$ is the average intensity and $N(4\pi)^{-1}$ is the mean-square intensity, $\sigma=(N-\alpha_{00}^2)^{1/2}/\alpha_{00}$ is the normalized "scatter" of the luminaire, in which there has been some recent interest in the illuminating engineering literature^{5, 6}. Therefore σ essentially measures the dominance (or lack thereof) of the first term of the expansion with respect to all of the others.

4 Simulated Demonstration and Finite Sampling

A simulated (simple, noise-free) source was used to demonstrate the rotational invariance of the d_l in practice. The illumination pattern of a glowing hollow cylinder of constant radiance was determined by raytracing for first a horizontal orientation and then a vertical one. Because of the lack of noise, any discrepancies in this example will quantify the distorting effect of finite, discrete sampling: the precise locations of sampled radiator patches will not exactly correspond between the two datasets, and thus the two datasets are not actually perfectly congruent under any rotation. The stepsizes used here were 4.5 degrees each in θ and ϕ .

Although the two sets of $\{\alpha_{lm}\}$ coefficients are very different, when the d_l are calculated as shown in table 2,⁶ the fact that the glowing cylinders are actually alike (to two decimal places) becomes obvious; the effect of sampling on this scale is seen in the third decimal.

raw coefficients and computed d-numbers					
coef	vertical	horizontal	d	vertical	horizontal
$\alpha_{0,0}$	+2.490	+2.494	d_0	0.983	0.983
$\alpha_{2,0}$	-0.443	+0.223	d_2	0.175	0.176
$\alpha_{2,2}$	+0.000	+0.547	d_4	0.043	0.041
$\alpha_{4,0}$	-0.109	-0.039	d_6	0.022	0.023
$\alpha_{4,2}$	+0.000	-0.083	d_8	0.022	0.021
$\alpha_{4,4}$	+0.000	-0.109	d_{10}	0.014	0.014
$\alpha_{16,0}$	-0.017	-0.001	d_{12}	0.012	0.010

Figure 2: the coefficients are different but the d-numbers are similar

⁶For this case, all of the odd l d_l are zero by cylindrical symmetry.

5 Application to Real Sources

Four different coiled-filament headlamps⁷ were investigated. They are a horizontal filament "spot quality" lamp surrounded by a smooth glass bulb; a vertical axial filament "high beam" lamp with a cylindrical bulb that becomes conical on top; a vertical filament axial "low beam" lamp with a cylindrical bulb and a light-blocking cap on top; and a "dual beam" lamp with two horizontal filaments for high and low beams, with a cap on top. The spot quality bulb is used in the manufacturer's reflector design process to evaluate reflector prototypes, to hand finish dies, etc. The other three bulbs are standard headlamp bulbs, with part numbers 9005, 9006, and 9004 respectively (see appendix IV for photographs).

In practice, the goniometer operates by making a full sweep in ϕ from 0 to 2π radians before incrementing θ , which ranges from $-\pi$ to $+\pi$ radians.⁸ Thus each point is sampled twice (once with positive θ and once with negative θ), except for points on the "dateline" great circle which are sampled three times, and the pole which is sampled as many times as there are ϕ values. Multiple readings can be averaged to reduce the effect of detector drift over the course of the data gathering.⁹ The many readings at the pole are used to obtain an estimate of the uncertainty or standard deviation of the measurements, which is then used to estimate the uncertainties of the α_{lm} and, by extension, the uncertainties of the d_l by the usual "propagation of error" techniques³. Typical values for the reduced σ of the intensity readings at the pole are in the range of ten to twenty percent of the pole's mean, showing the datasets were plagued by noise (mostly of mechanical origin, i.e., "wobble"). Knowing the uncertainties would be important, for example, for creating a fuzzy-logic quality control system to compare the d -number profile of a light source with a standard library of profiles. The uncertainty in the data, coupled with the effect of finite sampling (section 4), implies the d_l are reliable to one or one and a half decimal places.

⁷provided by a major automotive lighting supplier

⁸Due to physical limitations caused by the base of the bulb, the goniometer is unable to sample the entire range of θ . The intensity in this region should be zero, however, because it is in the shadow of the base.

⁹A baseline current was probably "bucked out" of the data when originally taken, but this value has unfortunately been lost. Calculating coefficients without this constant baseline results in a spuriously low value of α_{00} , but leaves the other α_{lm} unchanged. Information regarding d_0 is therefore irretrievably lost, but the true relative magnitudes of all of the higher l d_l are preserved.

The intensity dataset for the spot quality bulb is shown pictorially in figure 3, where the horizontal direction represents azimuth ϕ and the vertical direction represents elevation θ . For this dataset, the stepsizes were 3 grads¹⁰ in θ and 2 grads in ϕ . The dataset is shown as a sine-mapped spherical projection so that the pixel areas are proportional to the areas actually sampled over the sphere. The top of the dataset corresponds with the top of the bulb.

After fitting, when regenerated and interpolated at 1x1 grad steps, the processed data look like figure 4. Some features of this picture are worth mentioning. The two large dark spots are the ends of the horizontal filament, and the vertical posts holding the ends of the filament are also just visible. The "rings" are not artifacts; they appear clearly on real data taken with a finer grid, using steps of 1 grad each in θ and ϕ , as shown in figure 5. These dark rings are apparently due to "see-through" between the filament coils for viewing angles in ϕ that are nearly perpendicular to the filament.

¹⁰100 grads = 90 degrees = $\pi/2$ radians. The stepper motors on our goniometer are built with 400 steps per revolution, so our measurements are naturally reported in integral numbers of grads.

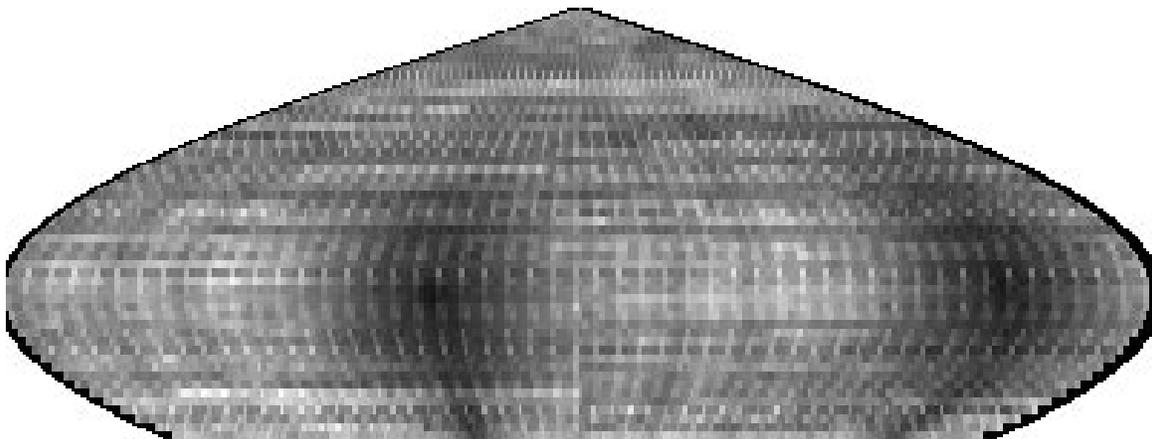


Figure 3: spot quality horizontal filament bulb, intensity data, 2x3 grad steps

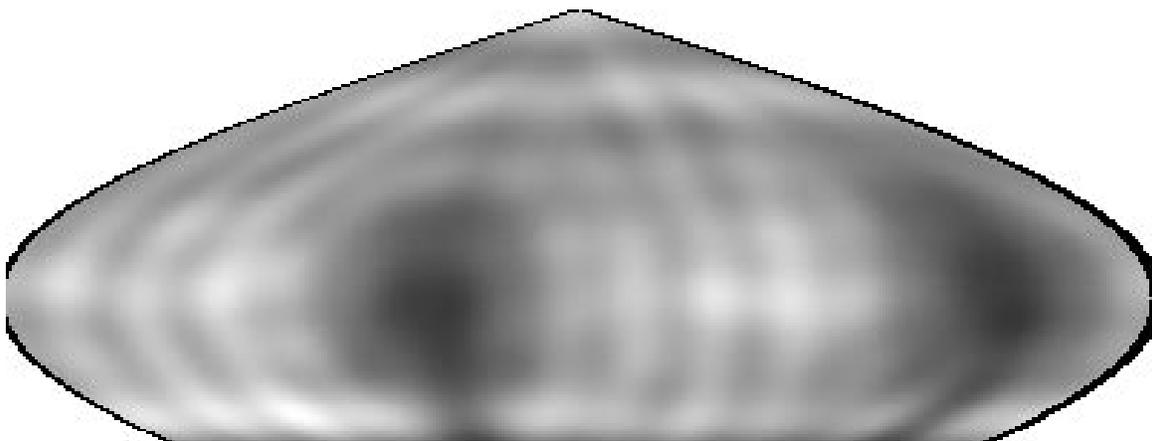


Figure 4: spot quality bulb, regenerated and interpolated data, 1x1 grad steps

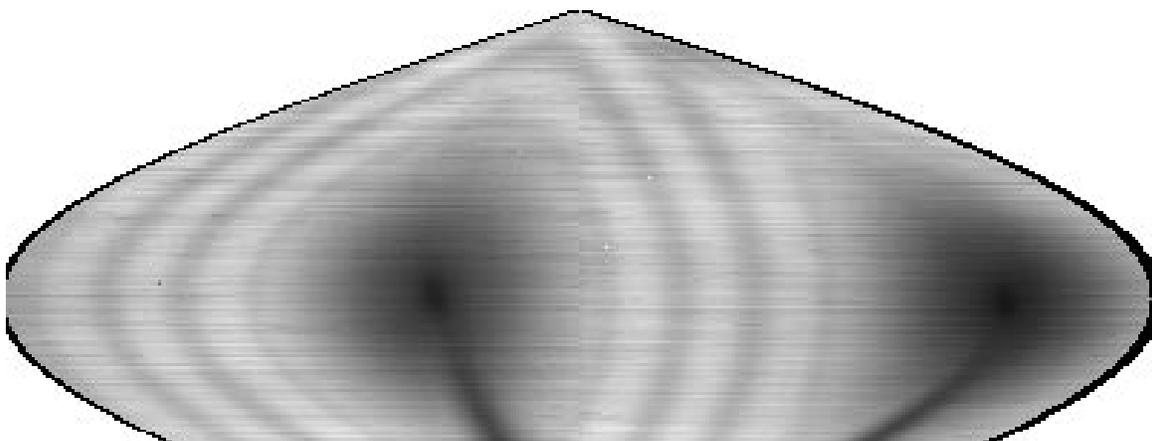


Figure 5: spot quality bulb, intensity data, 1x1 grad steps

A high beam lamp dataset is shown in figure 6, and its regeneration at six times the sampling density is shown in figure 7. Note the horizontal "see-through" rings near the center of the vertical filament. The dark stripe from pole to pole (of a few degrees azimuth) is the vertical filament mounting post; the other dark stripe 180° from the mounting post may be due to the post reflecting less light than the quartz bulb envelope.

Regenerated datasets for dual and low beam lamps are also shown in figures 8, 9, and 10.

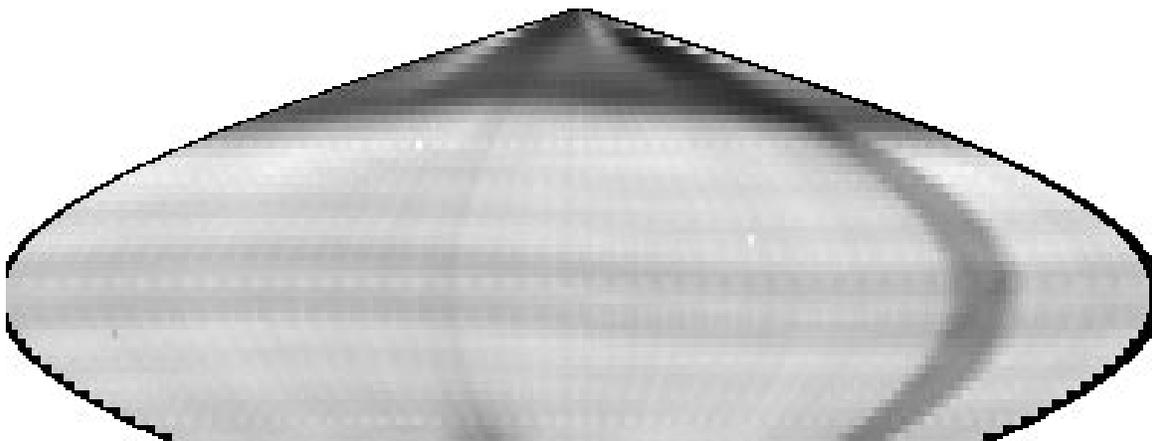


Figure 6: high beam bulb, intensity data, 2x3 grad steps



Figure 7: high beam bulb, regenerated and interpolated data, 1x1 grad steps

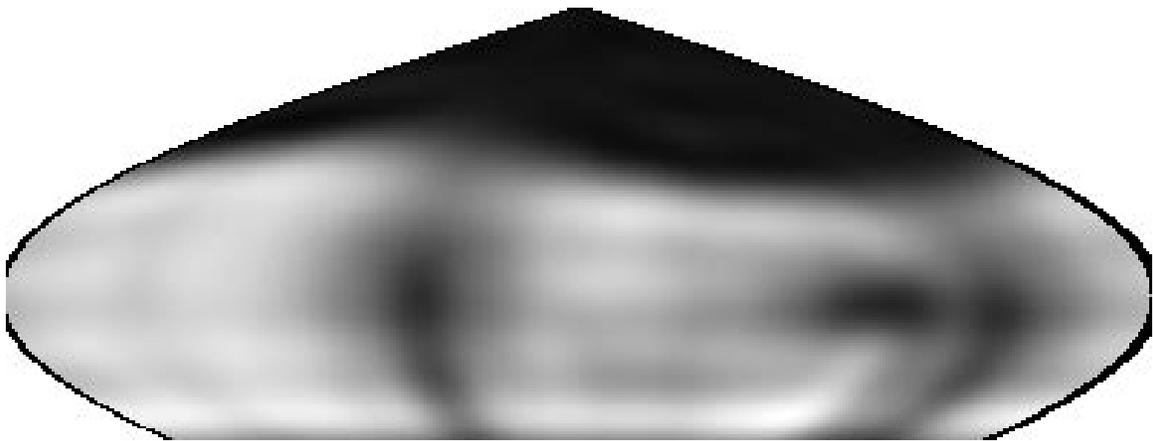


Figure 8: dual filament bulb, low beam, regenerated and interpolated data, 1x1 grad steps



Figure 9: dual filament bulb, high beam, regenerated and interpolated data, 1x1 grad steps

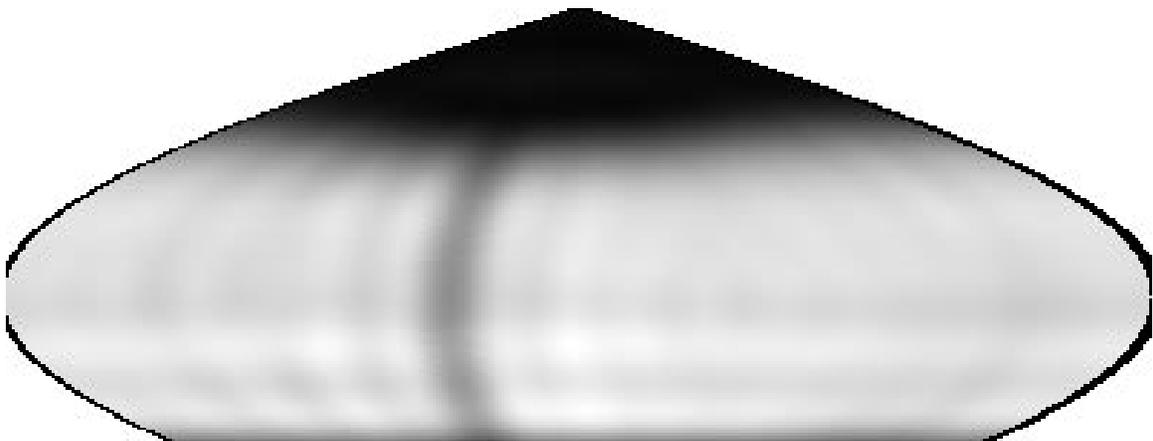


Figure 10: low beam bulb, regenerated and interpolated data, 1x1 grad steps

As expected, the intensity datasets for these filaments vary greatly, but the d -number profiles reveal certain similarities (to at least the first decimal) among the angular distributions of the various bulb types. The d -number profiles are shown in figure 11 for the first few values.¹¹ The high beam dual filament, for instance, produces an intensity dataset that is rotated by 180° about a vertical axis and has more power than the low beam dual filament, but their respective d -number profiles are essentially identical. The high beam and spot quality bulb are also similar due to the fact that both are identical filaments, but with differing orientations. Differences become apparent in higher orders, however, because the "holes" in the datasets, where data were unavailable, break the symmetries of the bulbs in different places. The low-beam bulb, with an axial filament and a cap to balance the shadow of the base, is nearly symmetric under parity¹² transformations and thus has the smallest odd- l d_l of all the bulbs.

A method for automated quality control in lamp production has been outlined in a paper by Lewin⁷. The typical quantities to be measured, however, are somewhat primitive, consisting of isocandela diagrams and average intensities; the d -numbers present a much more sophisticated standard for comparison.

6 Acknowledgements

This paper is submitted in partial fulfillment of the requirements for R. D. Stock's degree of Doctor of Philosophy at Carnegie Mellon University.

¹¹Values for d_0 are not shown for reasons stated in a previous footnote.

¹²An inversion of all coordinates.

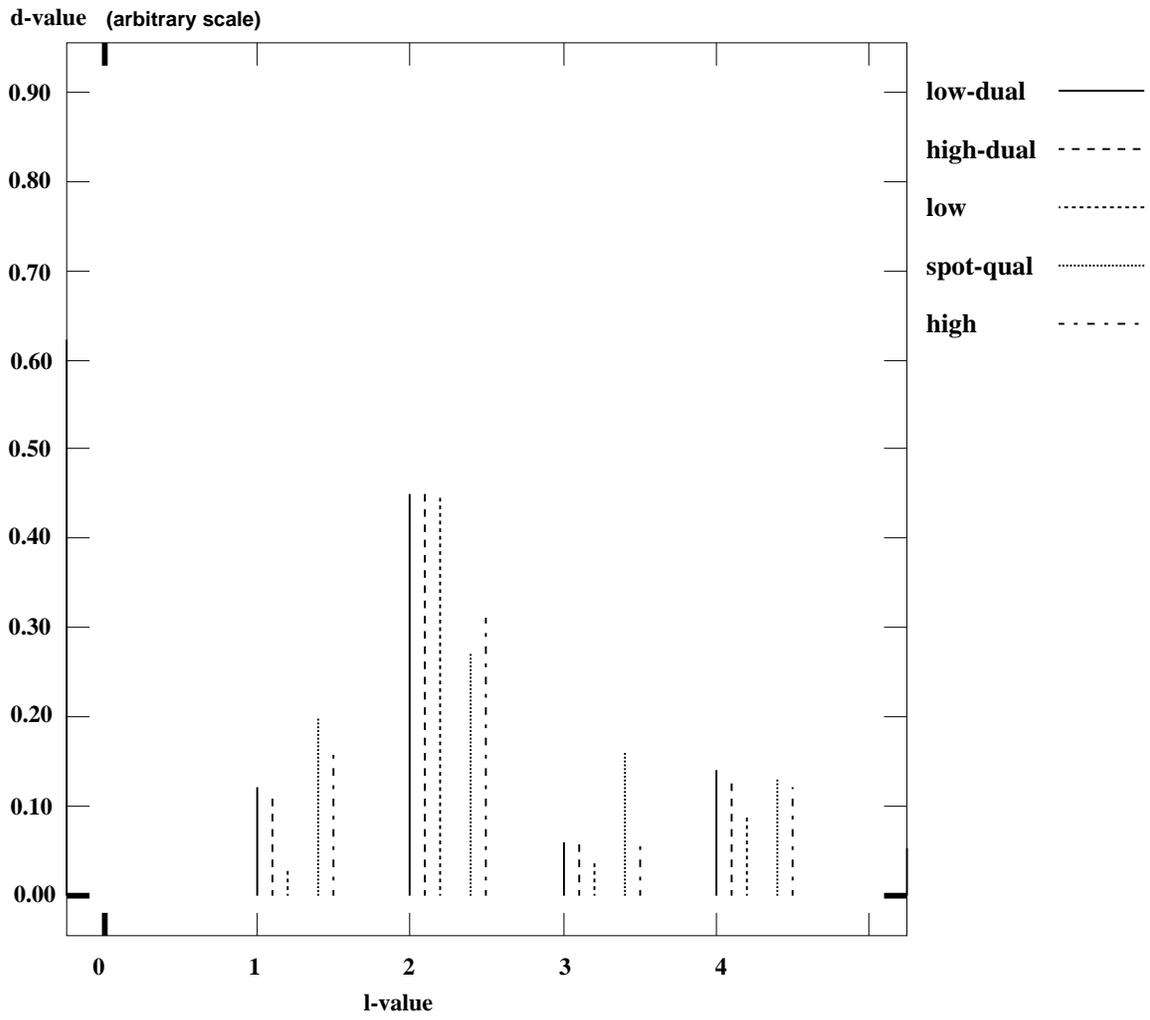


Figure 11: *d*-numbers for various bulb types

I. Proof of Rotational Invariance

Consider an arbitrary rotation of coordinates in which the point (θ, ϕ) is renamed (θ', ϕ') , resulting in a new set of coefficients $\{\alpha'_{lm}\}$. Requiring

$$I(\theta, \phi) = I(\theta', \phi')$$

results in

$$\sum_{lm} Y_{lm}(\theta, \phi) \alpha_{lm} = \sum_{lm} Y_{lm}(\theta', \phi') \alpha'_{lm}$$

Under a rotation of axes the Y_{lm} transform as⁴

$$Y_{lm}(\theta', \phi') = \sum_n D_{nm}^l Y_{ln}(\theta, \phi)$$

where the D are the Wigner D -matrices from Quantum Mechanics, and depend on the Euler angles of rotation of the axes. Substituting, we obtain

$$\sum_{lm} Y_{lm}(\theta, \phi) \alpha_{lm} = \sum_{lmn} D_{nm}^l Y_{ln}(\theta, \phi) \alpha'_{lm}$$

or

$$\sum_{lm} Y_{lm}(\theta, \phi) \alpha_{lm} = \sum_{ln} Y_{ln}(\theta, \phi) \cdot \sum_m D_{nm}^l \alpha'_{lm}$$

Because the spherical harmonics are orthonormal, i.e.

$$\int d\theta \int d\phi Y_{l'm'}^*(\theta, \phi) Y_{lm}(\theta, \phi) = \delta_{l'l} \delta_{m'm}$$

multiplying by $Y_{l'm'}^*(\theta, \phi)$ and integrating over angles produces

$$\alpha_{lm} = \sum_n D_{mn}^l \alpha'_{lm}$$

Therefore, for fixed l we have

$$\sum_m |\alpha_{lm}|^2 = \sum_m \left(\sum_p D_{pm}^{*l} \alpha'_{lp} \right) \cdot \left(\sum_q D_{mq}^l \alpha'_{lq} \right)$$

or

$$\sum_m |\alpha_{lm}|^2 = \sum_{pq} \alpha'_{lp} \alpha'_{lq} \left(\sum_m D_{pm}^{*l} D_{mq}^l \right)$$

The D -matrices are unitary, meaning

$$\sum_m D_{pm}^{*l} D_{mq}^l = \delta_{pq}$$

so that

$$\sum_m |\alpha_{lm}|^2 = \sum_m |\alpha'_{lm}|^2 = \text{constant under rotations.}$$

$$N \equiv \sum_{lm} |\alpha_{lm}|^2$$

and

$$d_l \equiv \left(\sum_m |\alpha_{lm}|^2 / N \right)^{1/2}$$

are also invariant under rotations. A set d_l depends on the shape of the angular distribution but not on the absolute intensity, which is related to α_{00} .

II. Data Fitting Details

The equation representing the integrated intensity dataset is

$$I[\Omega_i] = \sum_{lm} \alpha_{lm} Y_{lm}[\Omega_i]$$

where $[\Omega_i]$ is the i th measurement of $[\{\theta, \phi\}]$. This suggests the matrix equation

$$\mathbf{I} = \mathbf{a} \boldsymbol{\alpha}$$

Here

$$I_i = I[\Omega_i]$$

$$a_{ij} = Y_{lm}[\Omega_i]$$

and

$$\alpha_j = \alpha_{lm}$$

where the j in the previous two equations stands for the j th value of lm .

Because there are more measurements than coefficients, the system of equations can be solved by a least-squares ("LSQ") fit:

$$\mathbf{a}^T \mathbf{I} = \mathbf{a}^T \mathbf{a} \boldsymbol{\alpha} = \mathbf{b} \boldsymbol{\alpha}$$

Thus the coefficients are found by³

$$\boldsymbol{\alpha} = \mathbf{b}^{-1} \mathbf{I}$$

where the matrix \mathbf{b} is invertible by Gauss-Jordan elimination.

It should be noted, however, that although the LSQ method produces a low-residual fit, the coefficients it finds may be wildly "unphysical" due to an instability in the solution for the coefficients. This instability results from the fact that our measuring apparatus cannot sample the lower of the sphere surrounding the light source that contains the shadow of the base. The incompleteness of the data interval in θ causes a loss of orthogonality (implying a loss of linear independence) among the spherical harmonic basis functions, and consequently introduces an ambiguity into the contribution of some of the terms. Thus the values of the coefficients change as more terms are added⁸, and it becomes impossible to ascribe physical meaning to the d -numbers.

The point we wish to stress is that we are not simply interested in a good fit to the data, for which LSQ would be sufficient. Instead, we would like to extract information about the actual intensity distribution from the data. Some physical meaning must therefore apply to the

coefficients, and a necessary requirement is that the coefficients must be calculable incrementally, without new terms affecting the values of previous terms. Thus we found the LSQ method to be unsuitable for this calculation with this data; had the data interval been complete, the LSQ method would work (although a direct discrete Fourier dot product would be faster).

In theory, linear independence is a well-defined notion; two vectors are either linearly independent or they aren't. There is no room for dispute. In practice, however, especially in the realm of finite-precision numerical calculation, vectors can be partially dependent. In fact, it is possible to quantify mathematically the degree to which two vectors are linearly dependent^{9, 10}.

The method of singular value decomposition ("SVD") is specifically designed to address this problem^{9, 11, 8}. Although the matrix \mathbf{a} is not square, a pseudoinverse exists. Any arbitrary matrix can be written as the product of three trivially invertible submatrices:

$$\mathbf{a} = \mathbf{u}\mathbf{w}\mathbf{v}$$

If \mathbf{a} is m by n then so is \mathbf{u} , and \mathbf{w} and \mathbf{v} are both n by n . Both \mathbf{u} and \mathbf{v} are orthogonal, meaning the pseudoinverse equals the transpose. \mathbf{w} is diagonal, meaning its inverse is found simply by taking the reciprocal of its diagonal terms. Thus once this decomposition has been performed, the pseudoinverse of \mathbf{a} is found immediately to be

$$\mathbf{a}^{-1} = \mathbf{v}^T \mathbf{w}^{-1} \mathbf{u}^T$$

So far, nothing new has been gained. The diagonal terms of the matrix \mathbf{w} , called the "singular values" of the matrix \mathbf{a} , however, allow the identification of exactly which basis functions are becoming the most linearly dependent: setting the most ill-behaved w_{ii} to zero prevents those terms whose contribution to the Fourier sum is not uniquely determined from influencing the coefficients, whereas the LSQ method assigns large, delicately cancelling values to these coefficients¹⁰.

The relative stability of the two methods is shown in figure II-1, which gives the coefficient α_{00} as a function of l_{max} for each method. The solid line shows the LSQ method, and the dashed line shows the SVD method. The two methods essentially agree up to $l_{max}=4$, after which point SVD starts to zero out singular values. Thus a degree of stability can be achieved with SVD, but

the "black art" involved in deciding exactly how large a singular value must become in order to be considered "ill-behaved" makes this method still less than satisfactory.

One solution to this problem would be to use a new set of fitting functions that would be designed to be orthogonal over our specific data interval. The introduction of nonstandard functions, however, would mean a corresponding loss of both acceptance and physical intuition. Another approach to this problem would be to use an interpolating *sinc* or spline function. Because strict interpolations go through every data point, this method is not desirable for data that are noisy or that contain occasional large outliers. The idea is to smooth the noise over the entire dataset, rather than fit to it at each datapoint. For a given order, increasing the density of sampled points also reduces the ambiguity of the solution, but does not completely eliminate it.

The approach adopted, then, is to "fill in" the missing region with data of intensity zero, with the justification that this region corresponds roughly with the shadow of the base of the bulb. Because the spherical harmonic fitting functions are now orthogonal over our dataset, the least-squares fit becomes equivalent to a discrete Fourier "dot product" integral. Thus the coefficients are found directly by

$$\alpha_{lm} = \sum_i Y_{lm}^*[\theta, \phi]_i I[\theta, \phi]_i \sin \theta_i \Delta \theta \Delta \phi$$

where Y_{lm}^* is the complex conjugate of Y_{lm} , and the effort expended for an SVD matrix operation is avoided. The coefficients are also guaranteed to remain constant as more terms are added. Furthermore, if an LSQ matrix inversion is used with the complete filled-in dataset, statistical tools are built in³ that allow for the automatic computation of the uncertainties of the coefficients α_{lm} , although the computation is a little bit slower than the direct evaluation of the dot product.

Because the source is extended, the "shadow" region between light and dark should extend over a few values of θ . We have observed that smoothing the cutoff from light to dark linearly (for example), as opposed to creating a sharp drop, does indeed reduce the residual of the fit. Because the extent of the shadow region is not known, however, the results of this paper do not include smoothing of the cutoff.

An important remaining question concerns the maximum value of l used for the fit. The

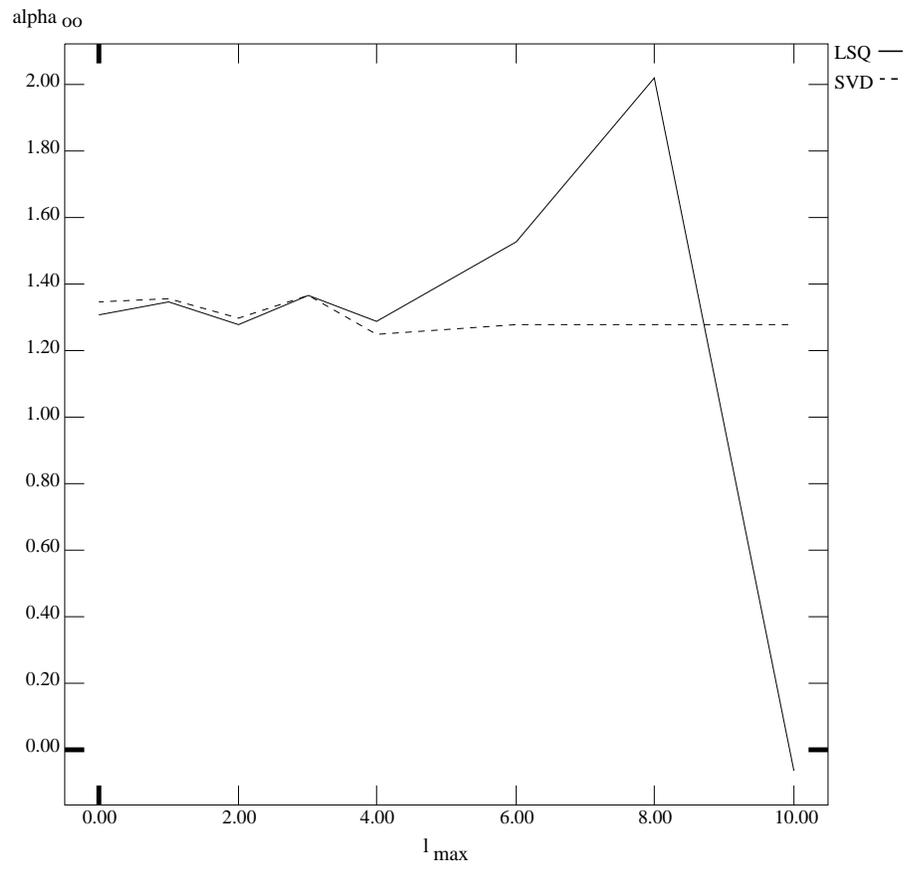


Figure II-1: stability comparison of SVD and LSQ solution methods

dataset pictures appear to regenerate well (by eye) typically around $l_{max}=16$; adding more terms does not seem to improve the fit significantly enough to justify the additional computation. A graph of the reduced χ^2 of the fit for a high beam bulb shows that $\chi^2 \rightarrow 1.2$ at $l_{max}=16$, beyond which it levels out (figure II-2).

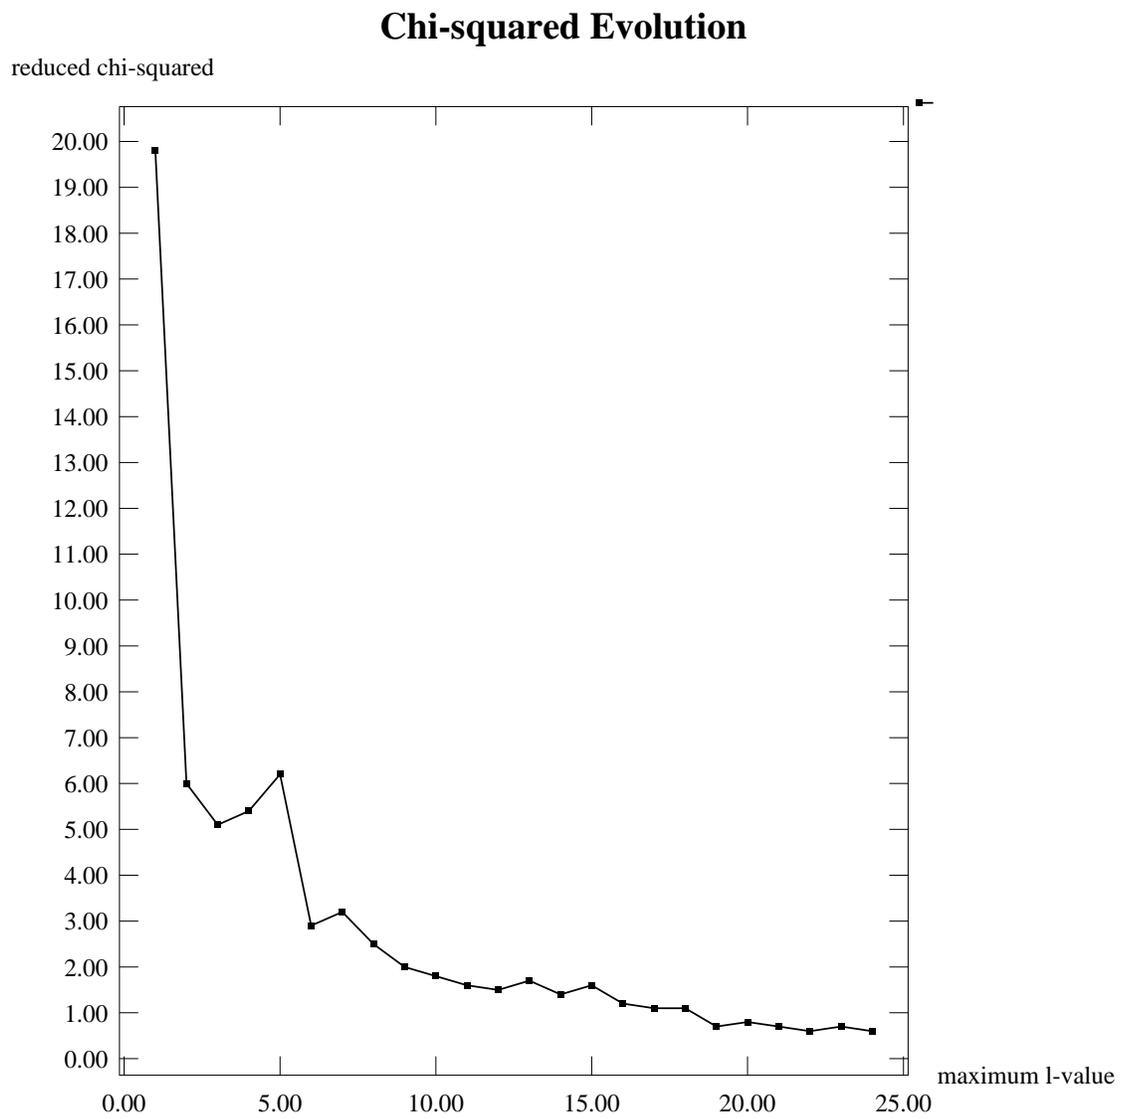


Figure II-2: graph of reduced chi-squared versus l_{max}

III. An Abstract Space

In an abstract, infinite-dimensional vector space spanned by an orthogonal set of unit vectors \mathbf{y}_{lm} , define the vectors

$$\mathbf{Y}(\theta, \phi) = \sum_{lm} Y_{lm}(\theta, \phi) \mathbf{y}_{lm}$$

and

$$\mathbf{A} = \sum_{lm} \alpha_{lm} \mathbf{y}_{lm}$$

Then the light source is represented in this space as the fixed vector \mathbf{A} , and the intensity at a point on the sphere surrounding the source is expressed as the dot product

$$I(\theta, \phi) = \mathbf{A} \cdot \mathbf{Y}(\theta, \phi)$$

Now define the subvectors

$$\mathbf{Y}_l(\theta, \phi) = \sum_m Y_{lm}(\theta, \phi) \mathbf{y}_{lm}$$

and

$$\mathbf{A}_l = \sum_m \alpha_{lm} \mathbf{y}_{lm}$$

In this representation, the d_l can be interpreted as the normalized length of the projection of the vector \mathbf{A} into the subspace l , or alternatively, as the normalized length of \mathbf{A}_l :

$$d_l = |\mathbf{A}_l| / |\mathbf{A}|$$

Note that

$$|\mathbf{A}| = \left\{ \sum_{lm} |\alpha_{lm}|^2 \right\}^{1/2} = \sqrt{N} = \sqrt{4\pi} I_{rms}$$

Physically, then, d_l^2 represents the rotationally invariant fraction of the total mean-square power distributed in the angular mode l .

Figure III-1 illustrates these relationships. Only three basis vectors (\mathbf{y}_{00} , \mathbf{y}_{10} and \mathbf{y}_{11}) are used for clarity. Regardless of the orientation of the measuring apparatus, the tip of the vector \mathbf{A} must lie on a circle of radius d_1 and height d_0 ; the closer \mathbf{A} is to vertical, the more uniform is the light source. Furthermore, valid vectors $\mathbf{Y}(\theta, \phi)$ representing points on the measuring sphere must also lie on a horizontal circle of fixed dimensions because the length

$$|\mathbf{Y}_l(\theta, \phi)| = (2l+1) / (4\pi)$$

is a constant independent of θ and ϕ .

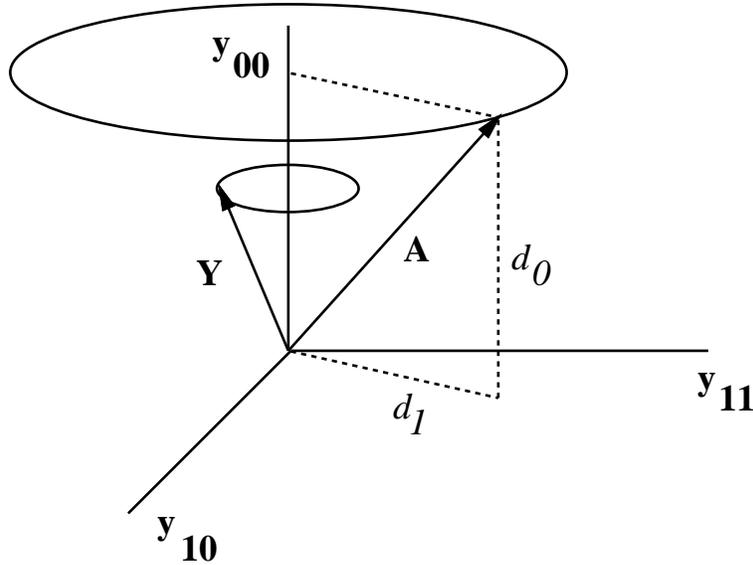


Figure III-1: vectors in an abstract space

In addition to providing an alternative description of the d_l , this representation also provides an amusing shortcut to the calculation of the coefficients α_{lm} in certain special cases. The dot product for $I(\theta, \phi)$ can be written as

$$I(\theta, \phi) = |\mathbf{A}| |\mathbf{Y}(\theta, \phi)| \cos(\beta)$$

Using

$$|\mathbf{A}| = \sqrt{4\pi} I_{rms}$$

and¹²

$$\begin{aligned} |\mathbf{Y}(\theta, \phi)| &= \left\{ \sum_{l=0}^{l_{max}} \sum_{m=-l}^{m=+l} |Y_{lm}(\theta, \phi)|^2 \right\}^{1/2} \\ &= \left\{ \sum_{l=0}^{l_{max}} (2l+1)/(4\pi) \right\}^{1/2} \\ &= (l_{max} + 1) / \sqrt{4\pi} \end{aligned}$$

results in

$$I(\theta, \phi) = I_{rms} (l_{max} + 1) \cos(\beta)$$

Therefore, in the (admittedly unlikely) event that the dataset happens to have the easily-tested property that

$$I_{max}(\theta_0, \phi_0) = I_{rms} (l_{max} + 1)$$

then $\cos(\beta) = 1$, which means \mathbf{A} is aligned with $\mathbf{Y}(\theta_0, \phi_0)$. Thus their coefficients must be directly proportional, i.e.

$$\alpha_{lm} = 4\pi Y_{lm}(\theta_0, \phi_0) I_{rms} / (l_{max} + 1)$$

or

$$\alpha_{lm} = 4\pi Y_{lm}(\theta_0, \phi_0) I_{rms}^2 / I_{max}$$

So for special distributions, the abstract space model gives the coefficients immediately, without the need for fitting the dataset to spherical harmonic functions.

IV. Bulb Types

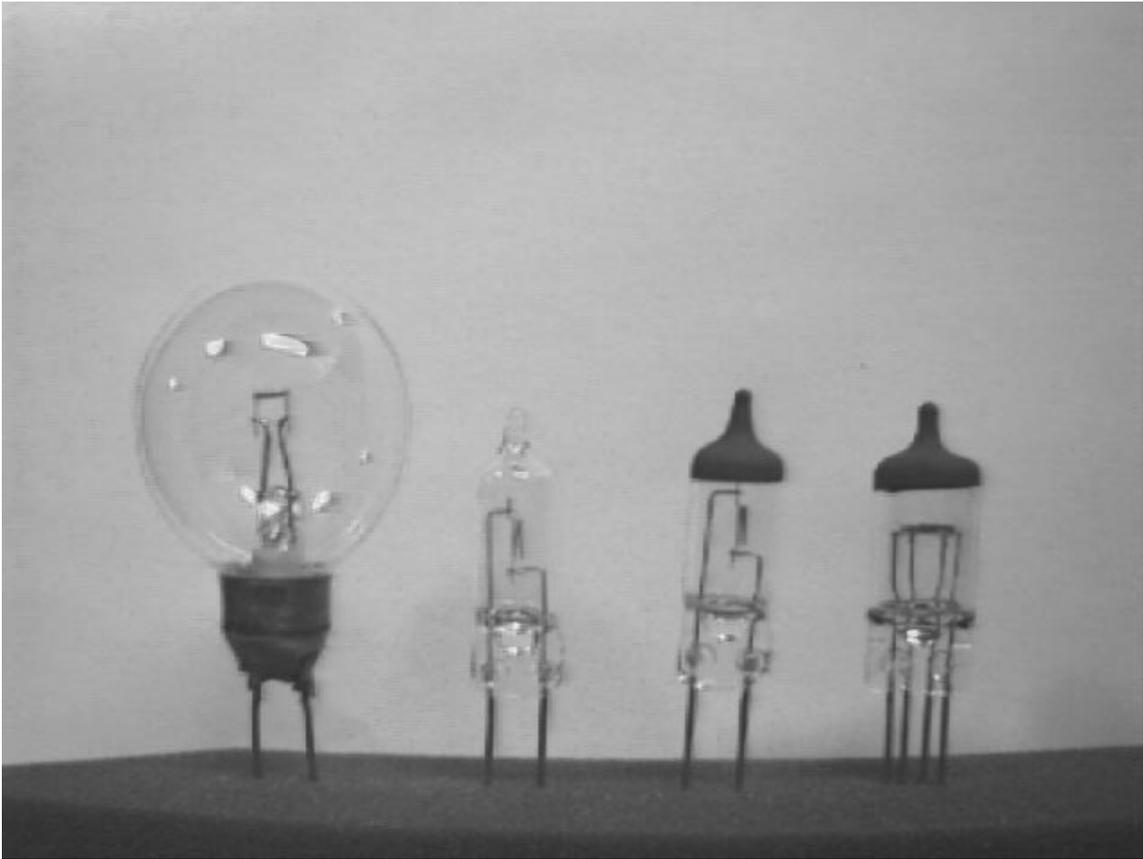


Figure IV-1: spot quality, high, low, and dual beam bulbs

References

1. C. Verbeck and D. Greenberg, "A Comprehensive Light-Source Description for Computer Graphics", *IEEE Computer Graphics and Applications*, Vol. 4, July 1984, pp. 66-75.
2. D. A. Varshalovich et al., *Quantum Theory of Angular Momentum*, World Scientific Publishing Co Pte Ltd, Singapore, 1988.
3. P. R. Bevington, *Data Reduction and Error Analysis for the Physical Sciences*, McGraw-Hill, New York, 1969.
4. M. E. Rose, *Elementary Theory of Angular Momentum*, Wiley, New York, 1957.
5. J. D. Armstrong, "A New Measure of Uniformity for Lighting Installations", *Journal of the Illuminating Engineering Society*, Vol. 19, No. 2, 1990, pp. 84-89.
6. J.-P. Mathieu, "Statistical Uniformity, A New Method Of Evaluation", *Journal of the Illuminating Engineering Society*, Vol. 18, No. 2, 1989, pp. 76-79.
7. I. Lewin, R. Laird, and B. Carruthers, "Development of New Photometer Concepts for Quality Control Applications", *Journal of the Illuminating Engineering Society*, Vol. 19, No. 2, 1990, pp. 90-97.
8. R. W. Hamming, *Numerical Methods for Scientists and Engineers*, McGraw-Hill, New York, 1973.
9. G. E. Forsythe, *Computer Methods for Mathematical Computations*, Prentice-Hall, Englewood Cliffs, NJ, 1977.
10. W. H. Press et al., *Numerical Recipes in C*, Cambridge University Press, Cambridge, 1988.
11. G. H. Golub and C. Reinsch, "Singular Value Decomposition and Least Squares Solutions", *Numerische Mathematik*, Vol. 14, 1970, pp. 403-420.
12. J. D. Jackson, *Classical Electrodynamics*, John Wiley & Sons, New York, 1975.

Table of Contents

1 Abstract	1
2 Data Collection and Analysis	1
3 Orientation Independent Source Characterization	3
4 Simulated Demonstration and Finite Sampling	5
5 Application to Real Sources	6
6 Acknowledgements	11
I. Proof of Rotational Invariance	13
II. Data Fitting Details	15
III. An Abstract Space	21
IV. Bulb Types	24

List of Figures

Figure 1: goniometer and detector apparatus	2
Figure 2: the coefficients are different but the d-numbers are similar	5
Figure 3: spot quality horizontal filament bulb, intensity data, 2x3 grad steps	8
Figure 4: spot quality bulb, regenerated and interpolated data, 1x1 grad steps	8
Figure 5: spot quality bulb, intensity data, 1x1 grad steps	8
Figure 6: high beam bulb, intensity data, 2x3 grad steps	9
Figure 7: high beam bulb, regenerated and interpolated data, 1x1 grad steps	9
Figure 8: dual filament bulb, low beam, regenerated and interpolated data, 1x1 grad steps	10
Figure 9: dual filament bulb, high beam, regenerated and interpolated data, 1x1 grad steps	10
Figure 10: low beam bulb, regenerated and interpolated data, 1x1 grad steps	10
Figure 11: <i>d</i> -numbers for various bulb types	12
Figure II-1: stability comparison of SVD and LSQ solution methods	18
Figure II-2: graph of reduced chi-squared versus l_{max}	20
Figure III-1: vectors in an abstract space	22
Figure IV-1: spot quality, high, low, and dual beam bulbs	24