# Maximum Entropy Lighting for Physical Objects

Tom Malzbender, Erik Ordentlich
HP Laboratories Palo Alto
HPL-2005-68
April 18, 2005*

image based lighting, graphics, information theory

This paper presents a principled method for choosing informative lighting directions for physical objects. An ensemble of images of an object or scene is captured, each with a known, predetermined lighting direction. Diffuse reflection functions are then estimated for each pixel across such an ensemble. Once these are estimated, the object or scene can be interactively relit as it would appear illuminated from an arbitrary lighting direction. We present two approaches for evaluating images as a function of lighting direction. The first uses image compressibility evaluated across a grid of samples in lighting space. The second uses image variance and prediction error variance, which are monotonically related to compressibility for Gaussian distributions. The advantage of the variance approach is that both image variance and prediction error variance can be analytically derived from the scene reflection functions, and evaluated at the rate of a few nanoseconds per lighting direction.

* Internal Accession Date Only
Approved for External Publication

# Maximum Entropy Lighting for Physical Objects

Tom Malzbender and Erik Ordentlich
Hewlett-Packard Laboratories

## Abstract

This paper presents a principled method for choosing informative lighting directions for physical objects. An ensemble of images of an object or scene is captured, each with a known, predetermined lighting direction. Diffuse reflection functions are then estimated for each pixel across such an ensemble. Once these are estimated, the object or scene can be interactively relit as it would appear illuminated from an arbitrary lighting direction. We present two approaches for evaluating images as a function of lighting direction. The first uses image compressibility evaluated across a grid of samples in lighting space. The second uses image variance and prediction error variance, which are monotonically related to compressibility for Gaussian distributions. The advantage of the variance approach is that both image variance and prediction error variance can be analytically derived from the scene reflection functions, and evaluated at the rate of a few nanoseconds per lighting direction.

## Introduction

Choosing good lighting directions is an important task for almost all human endeavors that involve vision, be it biological or computer vision. Here we present methods for choosing informative lighting directions to view real-world physical objects or scenes. These methods can also be applied to viewing synthetic objects generated by computer graphics, but are not limited to that domain.

The information conveyed to a human being by viewing an image or scene is of course a very difficult quantity to determine. Up to 50% of the human brain is involved in processing visual data in some manner, and we are far away from understanding this processing, especially that of 'late' or higher level vision. As a practical matter, therefore, in selecting informative lighting directions, we must rely on available proxies for the information content of an image. One natural proxy is the number of bits into which an image is encoded by a suitable data compression algorithm. Images requiring more bits to encode are taken to have greater information content. For example, an image in which all pixels are set to the same brightness level is encoded into only a few bits by virtually any reasonable compression algorithm, and, indeed, such an image would convey zero information to a viewer. In general, the true information content of an image, from a human being's perspective, and its compressed size can be widely divergent. An image consisting of pure noise, for example, is incompressible, while conveying zero information content. It seems reasonable to assume, however, that in suitably constrained sets of images, such as those arising in this work, true information content and compressibility, or, more accurately, incompressibility, are highly correlated. Our overall approach will, therefore, be to select lighting directions that result in the least compressible images.

As the basis for our compressibility measure, we consider two compression algorithms. The first algorithm is Huffman coding with respect to the empirical distribution of brightness values in the image. The overall encoding includes a description of the Huffman code used, as well as the encoding of the image data by the specified code. The compressed bit stream size for this algorithm, in bits per pixel, is given by (up to a negligible error term) the empirical entropy, H [Shannon 49]:

$$H = -\sum_{i=0}^{n} p_i \log p_i \qquad (1)$$

where I indexes the various brightness levels and $p_i$ is the fraction of times brightness $i$ occurs in the image.

First order image entropy, as measured by (1), is, of course, a limited measure of the information content. No spatial relationships are accounted for in this measure, nor are the relationships between discrete luminance bands. This motivates the second compression algorithm we consider: JPEG-LS, the current ISO standard algorithm for lossless image compression, which achieves superior compression, in part, by exploiting dependencies between neighboring pixel values [Weinberger 00] . It will be seen that in certain situations JPEG-LS compressibility can yield more "informative" lighting directions than first order entropy.

## Previous Work

There is a significant body of work dealing with the selection of good lighting directions. In the computer graphics literature most approaches have relied of user judgment to assess the quality of lighting choices [Marks 97], [Gershbein 00]. If one has a 3D model of the scene on hand, along with surface material properties, the approach of

[Shacked 01] can be used to determine optimal lighting parameters based on a number of perceptual metrics. An important advance is also presented in [Gumbold 02], where image entropy is used to automatically infer the quality of lighting. This model is extended to correspond more accurately with experimental data from human subjects evaluating the quality of lighting directions. Several limitation in Gumbolds approach are addressed in [Vazquez 03], namely the incorporation of color measures that may distinguish equi-luminance regions, and spatial continuity of patches of similar color and luminance. Vazquez also works in a perceptually uniform color space, namely LUV.

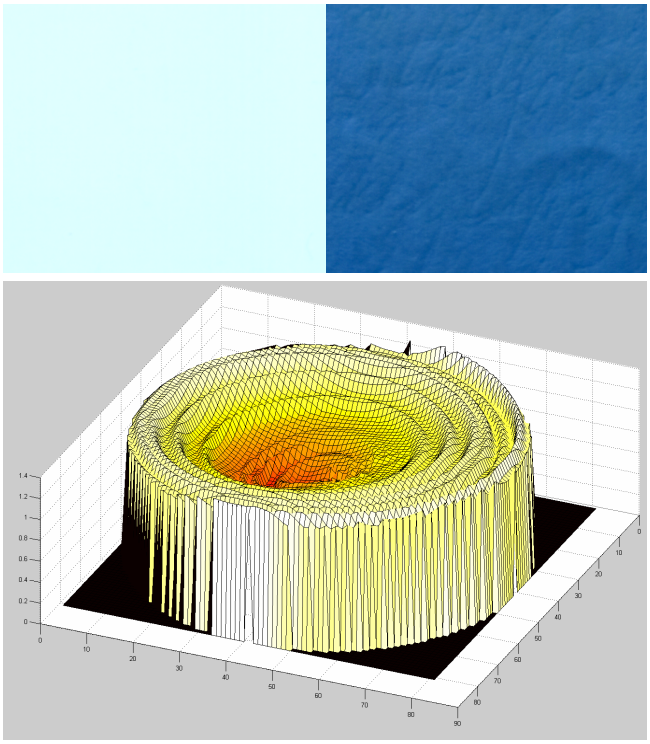

**Figure 1**. Top left: Sample of indented writing reconstructed from an overhead lighting position yielding low image entropy. Top right: Same sample using a grazing lighting direction (high entropy, low compressibility) yields image detail. Bottom: Entropy plot of writing sample as a function of lighting direction.

Like Gumbold and Shaked, our method is based on image entropy to measure the information content a particular lighting direction achieves. However, our approach is image-based and can be applied to any physical object, without requiring any 3D geometry. We start by taking 50 images of a static object or scene. In each image, illumination is provided by a point light source at a unique, known location. These measurements are then incorporated into a low order per-pixel reflectance model [Malzbender 01] of the form:

$$L(u,v;l_u,l_v) = a_0(u,v)l_u^{\,2} + a_1(u,v)l_v^{\,2} + a_2(u,v)l_u l_v + a_3(u,v)l_u + a_4(u,v)l_v + a_5(u,v) \qquad (2)$$

where $(l_u,l_v)$ are projections of the normalized light vector into the local texture coordinate system $(u,v)$ and $L$ is the resultant surface luminance at that coordinate. Coefficients $a_0$-$a_5$ are computed using least-squares fitting independently for each pixel from the 50 illumination varying images.

Once this characterization is complete, images can be computed from arbitrary lighting directions at greater than 60 frames / second using either graphics hardware acceleration or software only implementations, This is due largely to the reliance of eq. 2 on only fixed point multiplication and addition.

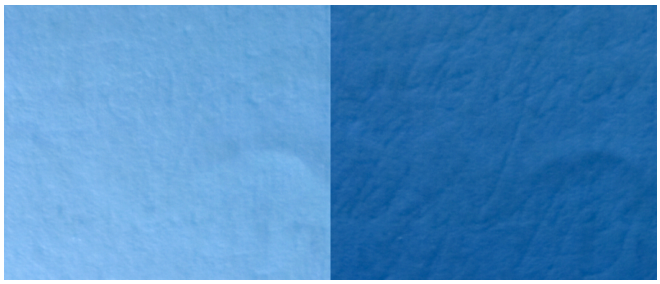

**Figure 2**. Images for lighting directions corresponding to maximum first order empirical entropy (left) and maximum JPEG-LS compressed size (right).

## Entropy Measure

To evaluate image compressibility, we uniformly sample the lighting space $(l_u, l_v)$ at discrete locations and compute the resultant image using PTM parameters $a_0$-$a_5$. From these images we can compute the histogram, and derive the image entropy; and we can also compress each image to derive its JPEG-LS compressed size. Figure 1 shows one example applied to indented writing. This paper sample was underneath a sheet of paper as it was written on and reveals very subtle depressions from the superimposed writing. When viewed from overhead lighting only the color of the paper is revealed, and measuring entropy and JPEG-LS compressed size yields low values (high compressibility) associated with a lack of detail. However, grazing lighting helps bring out features, consistent with increased entropy values and JPEG-LS compressed sizes (reduced compressibility) in the perimeter of the lighting space shown in Figure 1.

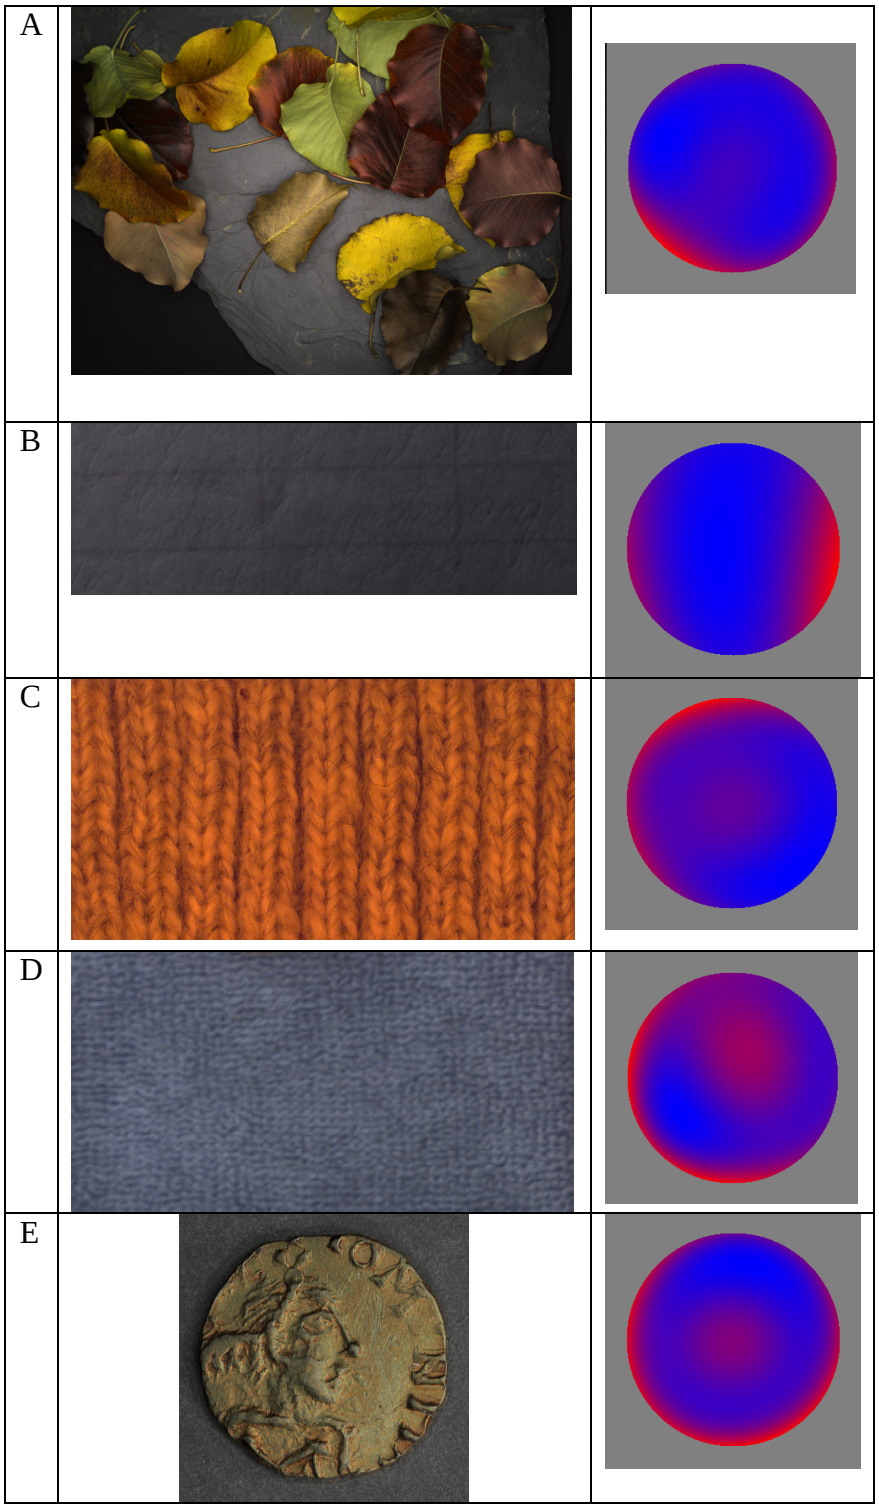


**Figure 3.** Objects captured under 50 varying lighting directions and associated image variance as a function of lighting direction. Red indicates lighting directions of high luminance variance, blue of low variance. A) Leaves on a slate background. B) Indented writing on paper. C) Knit sweater closeup D) Polyester carpet E) Gold Coin from the 7$^{th}$ century.

Figure 2 compares the images corresponding to maximum image entropy and maximum JPEG-LS compressed size, respectively, for the example of Figure 1. As can be observed, the least compressible JPEG-LS image is more informative than the maximum first order entropy image, in that the writing indentations are more visible. Apparently, in this case, the first order entropy measure is inflated by a gradual and systematic brightness variation across the maximum entropy image. The JPEG-LS algorithm, on the other hand, takes into account relationships between neighboring pixels and (roughly speaking) assigns shorter compressed codewords to effects (such as a systematic brightness variation) that can be learned and predicted as the image is scanned.

## Analytic Variance Functions

Evaluating the compressibility measures discussed above requires rendering an image at a particular lighting condition, then computing the compressibility measure from the image. This is also true of other measures such as those discussed earlier [Marks 97], [Gershbein 00], [Shacked 01], [Gumbold 02], [Vazquez 03]. However, in our image based approach using a low order parametric model (eq. 2), we derive an analytic description of how the surface reflectance varies with lighting direction.

Although the empirical first order entropy and JPEG-LS compressibility cannot be analytically derived from this representation, related quantities can. One such quantity is the empirical image variance. The link between variance and entropy is strongest for a Gaussian random variable for which the (differential) entropy is given by $(1/2)\log(2\bullet e\bullet \sigma^2)$, where $\bullet^2$ is the variance. We choose the following formulation for empirical variance where $L_i$ is pixel $i$'s luminance value in an image:

$$\sigma^2 = \overline{L_i^2} - \overline{L_i}^2 \qquad (3)$$

referring to Appendix A, we recover an analytic description of this reflection based image variance,

$$\sigma^2 = F(l_u, l_v) \qquad (4)$$

Once the 21 variance constants $(J_0\text{-}J_{14}, K_0\text{-}K_5)$ are computed from the parametric reflection coefficients, $a_0(u,v) - a_5(u,v)$, image variance from an arbitrary lighting direction can be evaluated at a cost of few dozen multiplies and adds (Appendix A). Images of several samples along with their empirical variance functions are shown in Figure 3. Red areas correspond to regions of high variance, typically corresponding to informative lighting directions. Note how grazing lighting is often preferred.

Similar analytical quantities can be derived for approximating JPEG-LS compressibility. One such quantity is the empirical variance of the prediction error corresponding to the linear portion of the predictor at the core of the JPEG-LS algorithm. The prediction error at location $u,v$ under consideration is

$$e(u,v) = L(u,v) - L(u-1,v) - L(u,v-1) + L(u-1,v-1)$$

and the quantity of interest is

$$\sigma_e^2 = \overline{e(u,v)^2}.$$

Note that the empirical mean of $e(u,v)$ will always be essentially zero, since the four terms comprising $e(u,v)$, summed over all the $u,v$, cancel, up to a negligible amount arising from the image boundaries. Since $e(u,v)$ is linear in the pixel values, the prediction error variance can also be expressed as a fourth order polynomial in $l_u$ and $l_v$, in a manner similar to (4), as it appears in Appendix A. In this case, additional correlations among the $a_i(u,v)$ are required. Let $\mathbf{a}(u,v)$ denote the vector $[a_0(u,v),\ldots,a_5(u,v)]$. The terms $J_0\text{-}J_{14}$ can be computed as linear combinations of the elements[1] of the 6 x 6 matrix $\bullet_{u,v}\ \mathbf{a}(u,v)^T\ \mathbf{a}(u,v)$. In addition to this, the prediction error variance also requires the matrices $\bullet_{u,v}\ \mathbf{a}(u,v)^T\ \mathbf{a}(u\text{-}1,v)$, $\bullet_{u,v}\ \mathbf{a}(u,v)^T\ \mathbf{a}(u,v\text{-}1)$, $\bullet_{u,v}\ \mathbf{a}(u,v)^T\ \mathbf{a}(u\text{-}1,v\text{-}1)$, and $\bullet_{u,v}\ \mathbf{a}(u,v)^T\ \mathbf{a}(u\text{+}1,v\text{-}1)$. The coefficients of the fourth order polynomial can then be computed as linear combinations of these matrix elements.

The prediction error variance, like its JPEG-LS compressibility counterpart, may be less susceptible to selecting suboptimal lighting directions than the first order empirical variance, particularly in cases similar to that of Figure 2. A simpler predictor resulting in the prediction error

$$e(u,v) = L(u,v) - L(u-1,v)$$

may also be considered, resulting in fewer correlations among the $a_i(u,v)$ required to compute $F(l_u,l_v)$, in this case only $\bullet_{u,v}\ \mathbf{a}(u,v)^T\ \mathbf{a}(u,v)$ and $\bullet_{u,v}\ \mathbf{a}(u,v)^T\ \mathbf{a}(u\text{-}1,v)$.

## Conclusions and Future Work

We have presented two methods for choosing lighting directions of physical objects. In both cases we start by collecting image-based per-pixel reflection measurements, which we parametrically describe using low order polynomial functions. In the first method, image compressibility is measured by constructing an image of the

[1] Note that this matrix is symmetric, so only 21 elements need to be computed.

object under a particular lighting direction and then compressing the resultant image. The second method involves a direct analytic derivation for image and prediction error variance, which can be computed directly from the reflection parameterization of the object and avoids the expensive image evaluation and compression/entropy measurement.

This second approach could also be applied across the space of possible enhancement parameters as well. For instance, [Malzbender 00] describes the technique of Diffuse Gain and its use in bringing out surface detail. The method has a free parameter, gain g, that controls the degree of enhancement one achieves. Although space does not permit its derivation here, an analytic expression for eq. 3 can be recovered, yielding a principled way to choose gain parameters *and* lighting directions that yield images with high variance.

## Acknowledgements

Thanks to Dan Gelb for collecting several of the reflectance functions shown in Figures 1,2, and 3. Thanks to Mark Mudge of Cultural Heritage Imaging for the use of reflectance functions of the gold coin appearing in Figure 3E. Figure 3E appears courtesy of the St. Bernard Pass monastery, Switzerland. Support for this project was provided by Fred Kitson, Mobile and Media Systems Lab.

## References

[Gershbein 00] Gershbein, R., Hanrahan, P., "A Fast Relighting Engine for Interactive Cinematic Lighting Design", Siggraph 2000, pp. 353-358.

[Gumhold 02] Gumhold, S., "Maximum Entropy Light Source Placement", Visualization 2002, Oct. 27-Nov. 1, 2002, Boston, Mass.

[Malzbender 01] Malzbender, T., Gelb, D., Wolters, H., "Polynomial Texture Maps", Siggraph 2001, pp. 519-528, Aug, 2001.

[Malzbender 00] Malzbender, T., Gelb, D., Wolters, H., Zuckerman, B., "Enhancement of Shape Perception by Surface Reflectance Transformation", HP Laboratories Technical Report, HPL-2000-38(R.1)

[Marks 97] Marks, J., Andalman, B., Beardsley, A., Freeman, W., Gibson, S., Hodgins, J., Kang, T., Mirtich, B., Pfister, H., Ruml, W., Seims, J., Shieber, S., "Design Galleries: A General Approach to Setting Parameters for Computer Graphics and Annimation", Siggraph 1997, pp.389-400.

[Shacked 01] Shacked, R., Lischinski, D., "Automatic Lighting Design using a Perceptual Quality Metric", Eurographics 2001, volume 20, pp. 215-226, Sept. 2001.

[Shannon 49] Shannon, C., Weaver, W., "The Mathematical Theory of Communication", University of Illinois Press, Urbana, Ill., 1949, ISBN 0-252-72548-4

[Vazquez 03] Vazquez, P., Sbert, M., "Perception-Based Illumination Information Measurement and Light Source Placement", Lecture Notes in Computer Science, 2003 (Proc. of ICCS'2003)

[Weinberger 00] Weinberger, M. J., Seroussi, G., Sapiro, G., "The LOCO-I Lossless Image Compression Algorithm: Principles and Standardization into JPEG-LS", *IEEE Trans. Image Processing,* Vol. 9, August 2000, pp.1309-1324.

## Appendix A
## Luminance Variance Derivation

Image variance, $\sigma^2$, can be analytically derived from the per-pixel reflectance functions described by eq. 2. Starting from the expression of variance for an image with *N* pixels described by eq. 3,

$$\sigma^2 = \overline{L_i^2} - \overline{L_i}^2$$

$$\sigma^2 = \sum_i \frac{L_i^2}{N} - \left( \frac{\sum_i L_i}{N} \right)^2$$

we substitute equation 2 yielding:

$$\sigma^2 = \frac{\sum_{u,v} (a_0 l_u^2 + a_1 l_v^2 + a_2 l_u l_v + a_3 l_u + a_4 l_v + a_5)^2}{N} - \left( \frac{l_u^2 \sum_{u,v} a_0 + l_v^2 \sum_{u,v} a_1 + l_u l_v \sum_{u,v} a_2 + l_u \sum_{u,v} a_3 + l_v \sum_{u,v} a_4 + \sum_{u,v} a_5}{N} \right)^2$$

simplifying yields:

$$
\begin{aligned}
&s^2 = F(l_u, l_v) \\
&= (J_0 + J_1 l_u + J_2 {l_u}^2 + J_3 {l_u}^3 + J_4 {l_u}^4 + J_5 l_v + \\
&J_6 l_u l_v + J_7 {l_u}^2 l_v + J_8 {l_u}^3 l_v + J_9 {l_v}^2 + J_{10} l_u {l_v}^2 + \\
&J_{11} {l_u}^2 {l_v}^2 + J_{12} {l_v}^3 + J_{13} l_u {l_v}^3 + J_{14} {l_v}^4) / N - \\
&\left( \frac{K_0 {l_u}^2 + K_1 {l_v}^2 + K_2 l_u l_v + K_3 l_u + K_4 l_v + K_5}{N} \right)^2
\end{aligned}
$$

where:

$$
J_0 = \sum_{u,v} {a_5}^2 ,\ J_1 = \sum_{u,v} 2a_3 a_5 ,\ J_2 = \sum_{u,v} {a_3}^2 + 2a_0 a_5 ,
$$

$$
J_3 = \sum_{u,v} 2a_0 a_3 ,\ J_4 = \sum_{u,v} {a_0}^2 ,\ J_5 = \sum_{u,v} 2a_4 a_5 ,
$$

$$
J_6 = \sum_{u,v} 2a_3 a_4 + 2a_2 a_5 ,\ J_7 = \sum_{u,v} 2a_2 a_3 + 2a_0 a_4 ,
$$

$$
J_8 = \sum_{u,v} 2a_0 a_2 ,\ J_9 = \sum_{u,v} {a_4}^2 + 2a_1 a_5 ,
$$

$$
J_{10} = \sum_{u,v} 2a_1 a_3 + 2a_2 a_4 ,\ J_{11} = \sum_{u,v} 2a_0 a_1 + {a_2}^2 ,
$$

$$
J_{12} = \sum_{u,v} 2a_1 a_4 ,\ J_{13} = \sum_{u,v} 2a_1 a_2 ,\ J_{14} = \sum_{u,v} {a_1}^2 ,
$$

$$
K_0 = \sum_{u,v} a_0 \ ,\ K_1 = \sum_{u,v} a_1 ,\ \ K_2 = \sum_{u,v} a_2 ,
$$

$$
K_3 = \sum_{u,v} a_3 ,\ \ K_4 = \sum_{u,v} a_4 ,\ K_5 = \sum_{u,v} a_5
$$