

IMPROVED SATELLITE BRIGHTNESS ESTIMATION TECHNIQUES

M.D. Hejduk

Titan Corporation

J.V. Lambert

Boeing Service Company

C.M. Williams

Volt Technical Services

R.L. Lambour

Massachusetts Institute of Technology Lincoln Laboratory

ABSTRACT

The Visual Magnitude Satellite Catalogue effort [1] presently employs a simple straight-line model for solar phase- and declination-angle dependencies, an approach that, at least for declination angle, lacked any real pedigree. The present attempt to improve the modeling began by investigating several competing gridding techniques that would allow two-parameter brightness dependencies to be displayed visually and selected Universal Kriging as the preferred methodology, as it can be demonstrated to produce minimum estimation error. Brightness histories from the SBV satellite were examined with krigged gridding for two difficult-to-model spacecraft: the HS-376 (cylinder) and HS-601 (cube with panels) geosynchronous satellites; and observed second-degree dependence on declination angle and, less commonly, phase angle suggested the use of a quadric surface fitting function. Quadric surface fits to the krigged gridding compared favorably to similar fits of the actual data while outperforming them in producing reasonable brightness extrapolation to areas of low data density. These results argue for a change of the VM SATCAT model to a quadric surface fit from krigged data, which, while not perhaps the final solution, would be a notable improvement over the present technique.

1. INTRODUCTION AND BACKGROUND

Time-intensity satellite brightness data have been used for many years for space object identification (SOI) and other applications. Discrete brightness estimates, generally a by-product of a metric observation and, when retained at all, used only at sensor sites and there only for mission planning purposes, are only now emerging as candidates for broader applications. The desire for SOI is increasing, and interest is growing in SOI-like products assembled from discrete brightness data. In support of the full-special-perturbations (SP) satellite catalogue, a new sensor tasking algorithm is under development, among whose needed inputs is accurate modeling of anticipated satellite brightnesses in order to take cognizance of sensor detectability in making tasking assignments. Finally, there is a requirement to characterize the brightness distribution of the satellite catalogue in order to make space debris estimates and define requirements for new sensors as the interest in tracking smaller objects grows. All of these needs argue for a robust satellite brightness model that can produce statistically-described brightness estimates for individual objects in a particular illumination geometry. Next year, it is expected that the GEODSS system will complete its DEEP STARE upgrade to CCD technology, and a feature of this upgrade is the retention of a well-calibrated visual magnitude measurement for every observation taken in sidereal mode; since this system of nine active cameras could well produce 30,000 brightness measurements in a favorable 24-hour period, the need for a robust model to represent this large data set is all the more acute.

The Visual Magnitude Satellite Catalogue effort, now at release 3.0, has been a first attempt to provide such a product. It uses as input brightness data from the Space-based Visible satellite (1996 to present), the Small-Aperture Telescope Augmentation experiment run at Edwards AFB (latter part of 1999), and the NASA Cloudcroft observing facility (1998-1999). Two features of this project that went somewhat beyond previous efforts were developing brightness models for as much of the satellite catalogue as reasonable data were present (about 90% of geosynchronous and 40% of DS satellites) and producing models that were not functions of only solar phase angle but also solar declination angle. The ability actively to fit most of the DS catalogue and achieve estimation errors generally less than a magnitude testified to the reasonableness of the approach, but it nonetheless suffers from many limitations. First, the model employed is very simple, presuming a relationship of the form

$$Mv = C_1 + C_2\alpha + C_3\delta . \quad (1)$$

Multiple studies [2,3,4], including a recent effort that attempted specifically to investigate behavior at large phase angles [5], have validated the linear relationship of satellite brightness to solar phase angle, at least for phase angles greater than 20 degrees; but there are no known study efforts that widely investigate the functional form of brightness dependency on solar declination angle. The linear assumption in (1) above seemed a reasonable initial guess but was based on rather little actual data examination. Second, no cognizance was taken of whether the satellite's orientation were actively maintained. The simple model might be adequate for debris or inactive payloads whose specular response is minimal and uncorrelated with viewing geometry; but objects whose stable payload orientation make their brightnesses subject to significant and predictable glint probably require something more elaborate. Finally, true random sampling is prevented both in phase angle, by sensors trying to optimize brightness (and therefore detectability) by scheduling as many of their observations as possible at low phase angles, and in declination angle, by sensor outages, seasonal changes in operating time, and the fixed time-period of experimental data preventing a temporally-uniform tracking volume spread over the whole of a year. This problem leads to the question of whether brightness data should be gridded before it is fit.

2. GRIDDING CONSIDERATIONS

Even highly-sampled satellites exhibit significant irregularities in the phase- and declination-angle distributions of their brightness samples. The data-density contour plot in Figure 1 below, typical of a highly-sampled satellite in the VM SATCAT, shows the notable regions of undersampling or complete absence of brightness data.

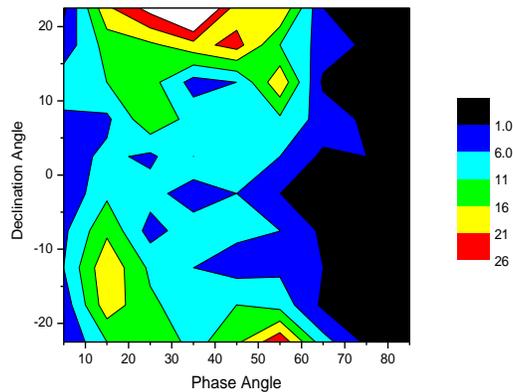


Figure 1: Data density contour plot for object 23199 (HS-376 cylinder). Despite the presence of 659 brightness samples, there are still significant regions of no data at all (black) or fewer than 6 samples per 10 x 5 degree frame (blue).

What is the best way to construct a full functional “map” of an attribute when the samples are random (in terms of independent variable values) and sparse? This problem has been thoroughly investigated in geology, in which, for example, topographic maps, water table contour plots, and ore-density diagrams frequently must be constructed from irregular and sparse data, often data that are available serendipitously (such as water table height data from wells that happen to exist) and cannot be gathered in accordance with design-of-experiments practices. Three basic gridding techniques are frequently encountered in geology, and it is helpful to explain them briefly to consider their relative merits.

A proximity moving average is the most basic of gridding approaches. To determine the value of each grid point, a circle (or sphere) is expanded until a set number of data points is encountered; and a weighted average of these points is calculated, with each point's weight determined by its distance from the grid point (usually $1/d$ or $1/d^2$). The technique's principal advantage—its ease of comprehension and implementation—is usually considered to be overshadowed by its disadvantages: it is heavy-handed in that it masks local detail very quickly, it performs poorly at the edges of the grid, and it does not ensure a good data sample about the grid point because all of the sample points might come from one “side” of the grid point. One correction for this latter problem is to require that a certain number of data points come from each quadrant or even octant about the grid point; but this approach, while ensuring a certain data diversity, usually requires that the point selection circle be expanded so broadly that all local detail is lost. In order to improve local accuracy, a second technique was developed that, rather than using a weighted average of the selected data points, fits a surface to the selected points and uses the surface to estimate the

grid point. This usually results in a better local fit, but it can produce wildly unrealistic values some distance from the data points used (*e.g.*, in trying to estimate values close to the grid perimeter or lacunae).

Kriging (named after D.G. Krige, a South African mining engineer) is a gridding technique developed to overcome these difficulties [6]. It makes use of regionalized variable theory, in which a variable is seen to vary continuously in any given neighborhood but not in a fully deterministic way; this is how many geological phenomena (and, it is postulated, satellite brightnesses) vary. To try to characterize a regionalized variable, one wishes to determine the degree of regional dependence as a function of distance from an arbitrarily-selected point. The building block of such a characterization is the *semivariance*, which, for a fixed sampling interval, is calculated as follows for a distance h from an arbitrary point:

$$\gamma_h = \sum_i^{n-h} (X_i - X_{i+h})^2 / 2n. \quad (2)$$

If this quantity is calculated for a large range of values of h , a plot of these data (and functional approximation of the plot) is called a semivariogram. Figure 2 below illustrates a canonical semivariogram and its core properties.

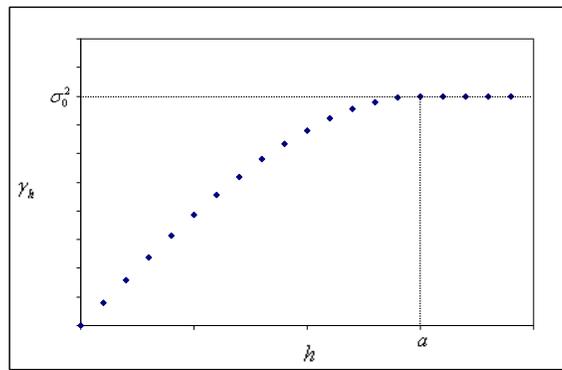


Figure 2: Canonical semivariogram conforming to spherical model. Sill is indicated by σ_0^2 , range by a .

As the graph makes clear, the semivariance γ_h increases with distance h until the increase tapers off and the semivariance becomes invariant with h . At this distance, called the *range* and denoted by a , the two points are no longer correlated at all; and their semivariance becomes equal to the overall variance σ_0^2 , called the *sill*. Ideal semivariograms pass through the origin (there is no variation with distance at a distance of 0) and reach a sill at an appropriate range. Semivariograms are first sketched empirically from sets of explicit calculations of (2) above; and when a durable dataset is assembled, the results are fit to one of several models. The canonical semivariogram model is the so-called spherical model, defined by the equation

$$\gamma_h = \sigma_0^2 \left(\frac{3h}{2a} - \frac{h^3}{2a^3} \right) \quad (3)$$

and used to create the example in Figure 2 above. When h reaches the range, the semivariance is set equal to the sill. Logarithmic and linear semivariogram models are also commonly used; all require some alteration to force the function to a sill value at some range.

If the correlation among data points as a function of their relative distance is known, this information can be used to calculate weighted averages for proximate average gridding in a manner that minimizes the estimation error. The full derivation of the minimization equations will not be provided here [7], but it follows the usual procedure of beginning with the equation for the error variance, substituting a random variable function for the unknown “true” values of the estimated points, taking partial derivatives with respect to the weighting variables desired, and setting these partial derivatives equal to zero as an expression of relative minima. For the case in which a grid point were to be estimated from three known data points, the minimization equations would be the following:

$$\begin{aligned}
w_1\gamma_{h_{11}} + w_2\gamma_{h_{12}} + w_3\gamma_{h_{13}} &= \gamma_{h_{1p}} \\
w_1\gamma_{h_{21}} + w_2\gamma_{h_{22}} + w_3\gamma_{h_{23}} &= \gamma_{h_{2p}} , \\
w_1\gamma_{h_{31}} + w_2\gamma_{h_{32}} + w_3\gamma_{h_{33}} &= \gamma_{h_{3p}}
\end{aligned} \tag{4}$$

in which w_x is the averaging weight to assign to point x and $\gamma_{h_{xy}}$ is the semivariance between points x and y (the required semivariances can be determined from the developed semivariogram for the dataset). Requiring that the estimate be unbiased introduces a fourth constraint:

$$w_1 + w_2 + w_3 = 0 , \tag{5}$$

making the problem overdetermined; this bit of good fortune can be used to introduce a Lagrange multiplier to force a minimum-error solution. The final set of equations to be solved thus takes the form

$$\begin{aligned}
w_1\gamma_{h_{11}} + w_2\gamma_{h_{12}} + w_3\gamma_{h_{13}} + \lambda &= \gamma_{h_{1p}} \\
w_1\gamma_{h_{21}} + w_2\gamma_{h_{22}} + w_3\gamma_{h_{23}} + \lambda &= \gamma_{h_{2p}} , \\
w_1\gamma_{h_{31}} + w_2\gamma_{h_{32}} + w_3\gamma_{h_{33}} + \lambda &= \gamma_{h_{3p}} \\
w_1 + w_2 + w_3 &= 0
\end{aligned} \tag{6}$$

which is four equations in four unknowns. Once solved, the estimated value for the grid point p is given by

$$\hat{p} = w_1p_1 + w_2p_2 + w_3p_3 ; \tag{7}$$

and the error variance for that estimated value, which is the weighted sum of the semivariances from each of the data points to the grid point, is determined by

$$s_e^2 = w_1\gamma_{h_{1p}} + w_2\gamma_{h_{2p}} + w_3\gamma_{h_{3p}} \tag{8}$$

The standard error of the estimate follows immediately by taking the square root.

Kriging advances itself as the preferred gridding solution because it both produces a solution with minimum estimation error and allows the calculation of an error variance for each grid point. This technique is not, however, without its own difficulties. If the entire two-dimensional surface to be gridded includes a slow change in average value across the entire surface, the surface exhibits “drift”, which must be removed before kriging and then added back after the kriging solution is obtained. Drift models are usually either linear or quadratic in character, and care must be exercised in model choice. Proper selection of the semivariogram model is also necessary for a reliable kriging solution. Finally, the number of points to use for each grid point estimate needs to be determined, and different sample size choices can alter the results. A kriging solution ready for “prime time” requires at least mild investigation of each of these items to ensure that judicious selections have been posited; and before any widespread use of kriging is introduced to the VM SATCAT effort, such investigations will take place. At the proof-of-concept level, however, it is sufficient to accept algorithmic defaults.

3. BRIGHTNESS BEHAVIOR INVESTIGATION: GENERAL

Satellite brightness behavior against solar phase and solar declination angles was investigated by dint of surface plots of these data, generated by kriging techniques. The kriging routine available with the analysis package of Origin 6.1 was used, with a data limitation of 0 – 90 degrees for solar phase angle (which excluded very little data, as there were only a few observations taken at an angle greater than 90 degrees) and set to form an even 32x32 grid (to prepare the gridded data for future Fourier transformation or wavelet decomposition, if desired). The routine was

unfortunately unstable, failing in about 25% of the cases due to sorting and matrix inversion errors. Because the source code for the interpreted routine was not available, few remediation avenues presented themselves. Once this study was complete, it was discovered that a MATLAB-compatible kriging product has been developed by the Woods Hole Oceanographic Institution and is freely available to engineers and scientists; this package is a promising candidate to replace the Origin routine should full-catalogue kriging be desired. Fortunately, the number of successful cases in this study was large enough to allow preliminary conclusions to be drawn, as there were no common characteristics among the cases that failed and thus no reason to suppose that the exclusion of the failures introduced a bias.

The surface plots were examined visually to determine what sort of model fitting function they suggested. This sort of evaluation is heuristic and bounded by the mandate to generate a simple model that can fit most cases reliably. While there were occasional suggestions of cubic or logarithmic behavior, in general it appeared that a quadric (second-degree) surface would work reasonably well, if degenerate cases were permitted. A second-degree function also avoids the instabilities of higher-degree polynomials and singularities associated with logarithmic functions. The general form of the proposed fitting function is thus

$$A\alpha^2 + B\alpha\delta + C\delta^2 + D\alpha + E\delta + F = z , \quad (9)$$

in which alpha represents solar phase angle and delta represents solar declination angle. The including of the rotation term(s) has been shown to be prudent in examining the plots. The fits were accomplished by the usual least-squares regression technique. An F -test was applied to the regression results, and the fit was in principle rejected if a significance value of 0.05 could not be achieved (this turned out to be a superfluous condition, as all attempted fits achieved at least this level). A t -test was then applied to each regressor variable, and a variable was rejected as not contributing significantly to the regression if a t -test significance level of 0.05 could not be achieved. In such a case, the failing variables were eliminated and the regression re-run with the significant regressor variables only; this mitigates against colinearity and ensures more meaningful r^2 values. Each results set includes a bar graph of the number of fits for which each of the regressor variable types contributed significantly to the regression; it is interesting that solar declination angle is at the least on par with phase angle as a significant contributor to the regression and often displaces it for this distinction.

The fits were executed against both the raw brightness data and the krigged grid data. In general, these two fit approaches did not yield significantly different results in that they both produced fits with the same “look and feel.” The fits against the krigged data did, however, outperform the raw-data fits in two respects. First, they yielded more realistic predictions of expected performance at the extremes of the grid space. Admittedly, this type of evaluation is circular in that it introduces an *a priori* expectation of such behavior even though the study is supposed to be determining what that behavior actually is. However, fits that show substantial changes in predicted brightness away from the more dense data areas, especially if the change in direction is counterintuitive, can be considered inferior to those that behave in a more moderate and continuous way. Second, for each fit a modified error statistic is calculated: the amount the constant term of the fit would need to be increased in order for a given percentage of the data points to fall *below* (i.e., brighter than) the fit surface. One of the principal uses for the VM SATCAT is to aid sensor tasking assignments; and in assigning tasking, one wishes to achieve a certain confidence that for the postulated viewing geometry the satellite’s brightness will exceed a given minimum value. This error statistic is computed for both fit types (to raw data, evaluated against the data points; and to krigged gridding, evaluated against the grid points), and in general the krigged fits outperform the data fits, often by close to a factor of two. Because the krigged situation will reproduce all the data points exactly (which is why it is referred to as an “exact interpolator”), these two sets of error statistics can be compared. These results will be shown in the ensuing sections, as the specific results from the two examined bus types are presented.

4. BRIGHTNESS BEHAVIOR INVESTIGATION: HS-601

The HS-601 bus is a traditional rectangular-prism, 3-axis-stabilized satellite bus, powered by deployable solar arrays and used to host a variety of communications payloads. A picture is given in Figure 3. The plethora of protrusions (solar panels, communications antennae, &c. predispose the satellite to frequent and significant glint returns. All HS-601 objects investigated here were actively stabilized according to the open-source *Satellite Encyclopedia*.

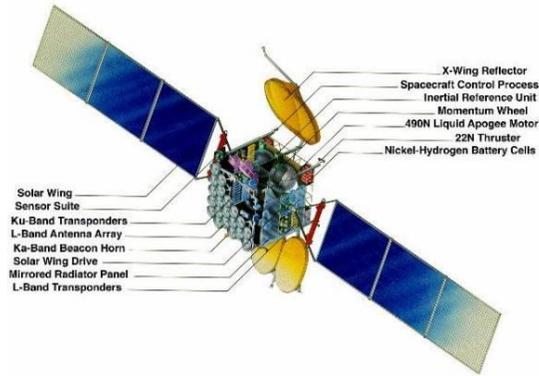


Figure 3: HS-601 communications satellite. Figure courtesy *Encyclopedia Astronautica*.

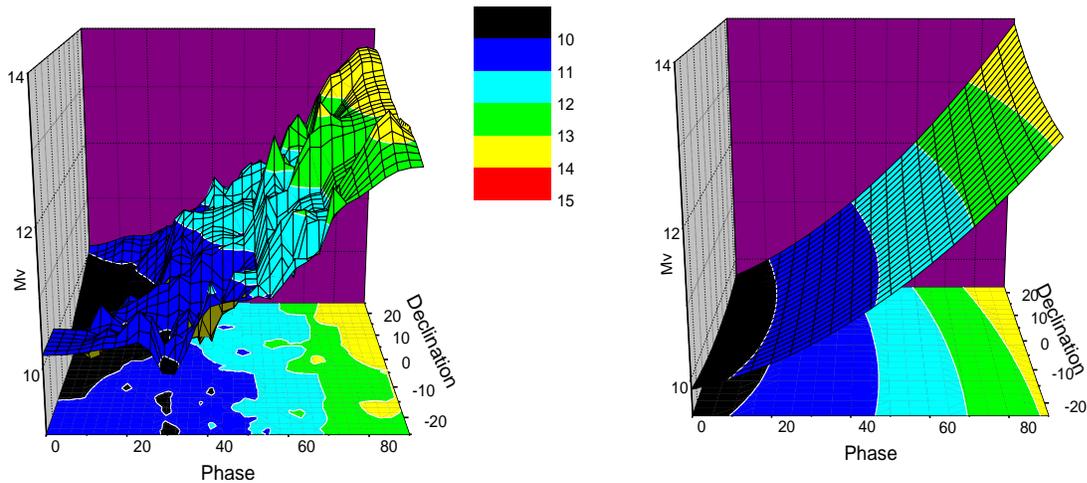


Figure 4: Krigged surface plot and second-degree fit for object 23132. The contour lines' near parallelism to the declination axis indicates their independence from declination.

About two-thirds of these cases manifested linear or near-linear brightness variation against solar phase angle and very little variation with solar declination angle, and indeed some second-degree fits to the data eliminated all regressor variables except that of solar phase angle. This was hardly the majority case, however; more commonly at least one of the solar declination regressor variables was retained, and quite often both. A typical case manifesting weak declination-angle dependence, and the associated fit, are given in Figure 4. More declination dependence is observed in the remaining third, of which object 23313 is a good example. Figure 5 gives both the krigged and quadric surfaces. The “target” or “rainbow” quality of the contour plot characterizes these cases.

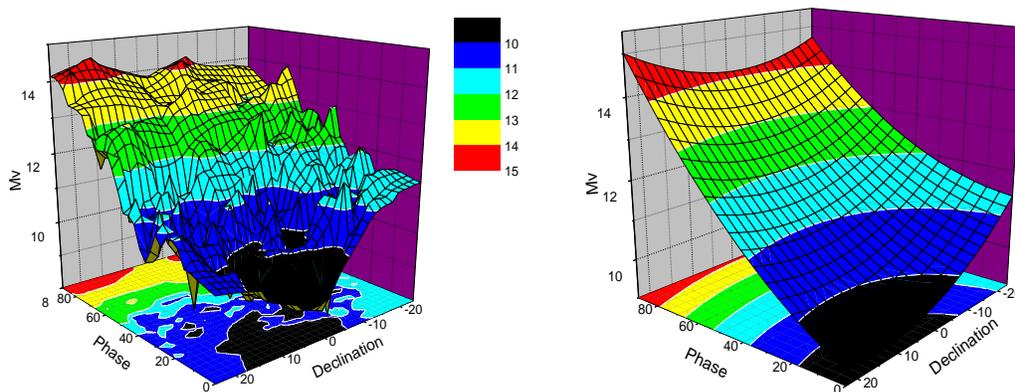


Figure 5: Krigged surface plot and second-degree surface fit for object 23313

Figure 6 gives some fit performance information. The left-hand graph gives the percentage of the fits in which each regressor variable contributed significantly to the regression. The fits to the “smoothed” krigged surface produce more elaborate quadric surfaces, but this is not particularly surprising. What does come as a surprise is that phase angle does not at all dominate the fitting; contribution is more or less even among the five variable types. The right-hand graph plots the error statistic described in Section 3 above for three percentile levels. The fits perform well here; an elevation of the fit surface of only 0.6 magnitude is sufficient to ensure that the surface serve as a 95th percentile upper bound.

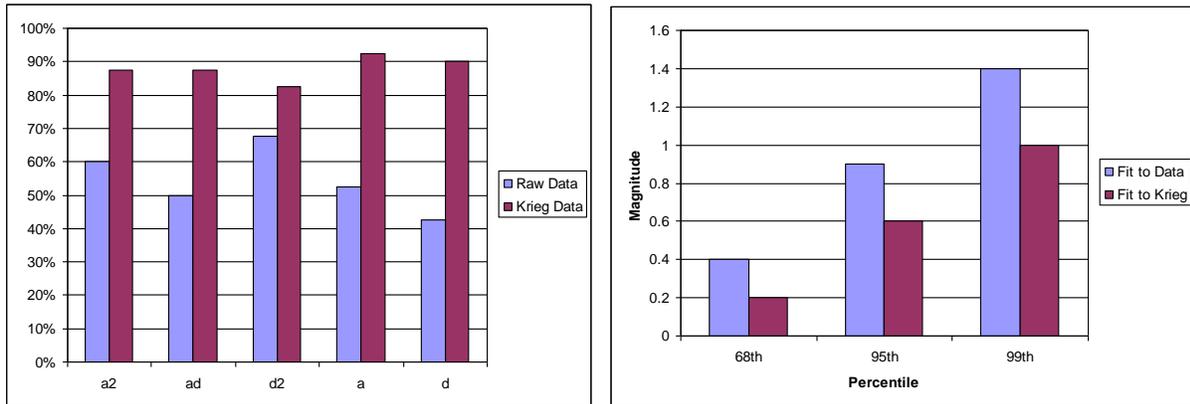


Figure 6: Regressor variables used and error statistics for HS-601 fits

5. BRIGHTNESS BEHAVIOR INVESTIGATION: HS-376

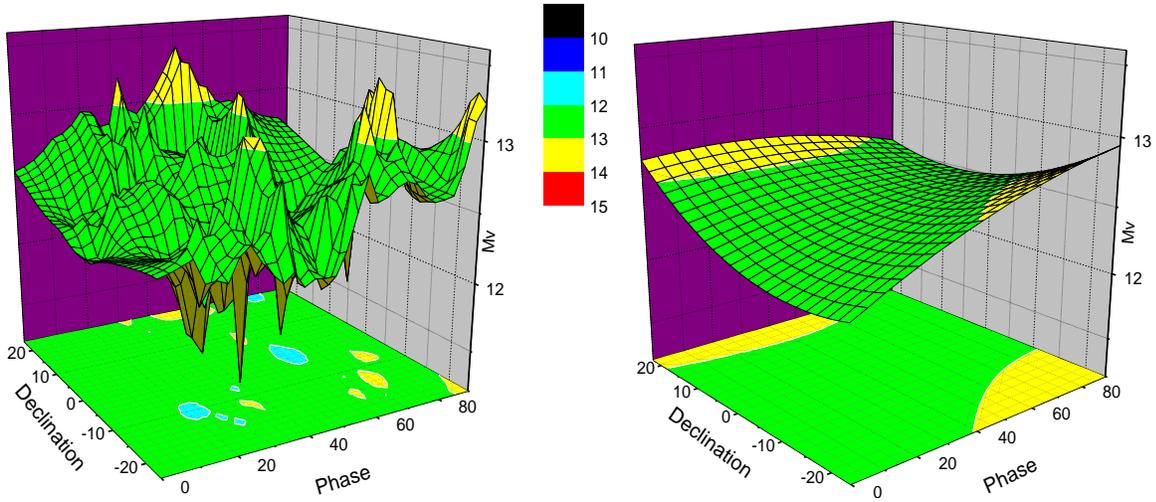
The HS-376 is a cylindrical, spin-stabilized communications satellite. Solar arrays are on-body, eliminating the need for active array management; and antennae and other communications apparatus reside at one end of the satellite that is de-spun and thus, in a geosynchronous orbit, remains in a fixed position relative to the earth’s surface. A picture of this spacecraft is given in Figure 7.



Figure 7: Generic version of HS-376. Figure courtesy *Encyclopedia Astronautica*.

Very different response was observed depending on whether the satellite were still actively stabilized, the source for which was the open-source *Satellite Encyclopedia*. Non-stabilized payloads’ brightness response tended to fall into two categories: very little dependence on any aspect of viewing geometry or a palpable but not overwhelming dependence on phase angle. The former is illustrated in Figure 8 (object 14158), which, although fitting both declination and the square of declination, produces such small regressor constants that the fit is really not much better than taking the mean ($r^2 = .08$). The latter is seen in Figure 9 (object 15993), in which phase angle influence dominates the fit. In both cases, the total brightness span is narrow, remaining between two magnitudes.

The fit results for the unstabilized case, summarized in Figure 10, show a surprisingly high contribution from declination and low contribution from phase, although it must be remembered that the correlations here are weak—the median r^2 value is only 0.2. The error statistics are pleasantly narrow; but in considering how narrow the actual data spread was, they do not seem all that impressive.



Figures 8 and 9: Krigged and fit data for satellite 14158 (above) and 15993 (below)

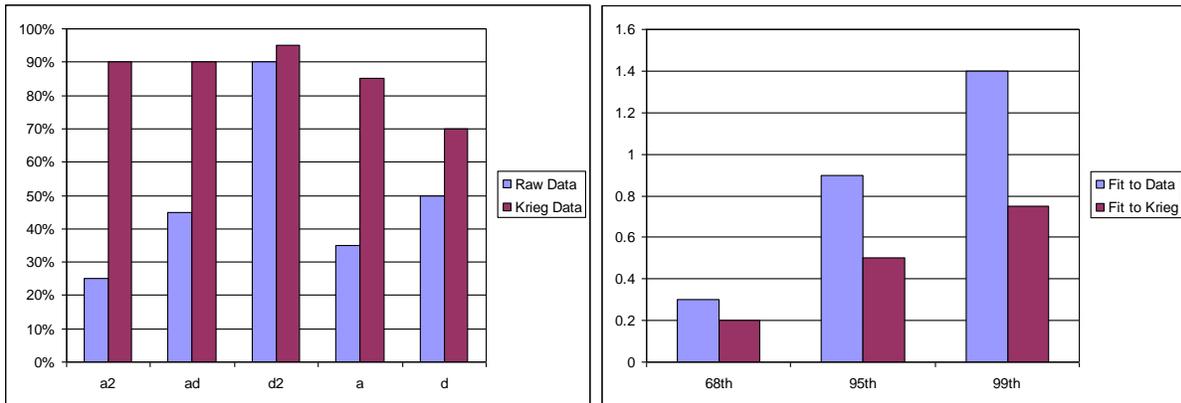
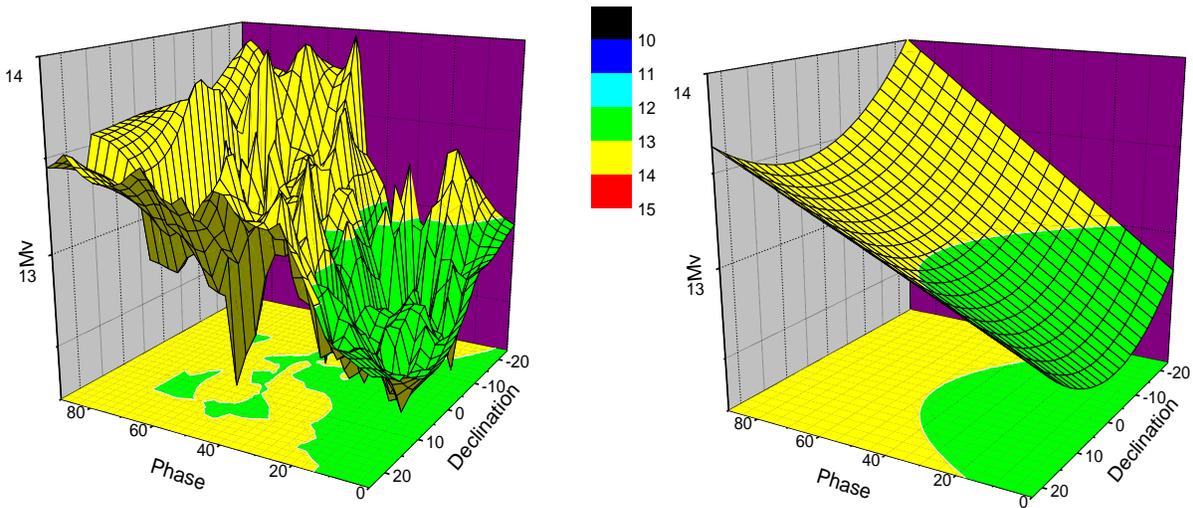
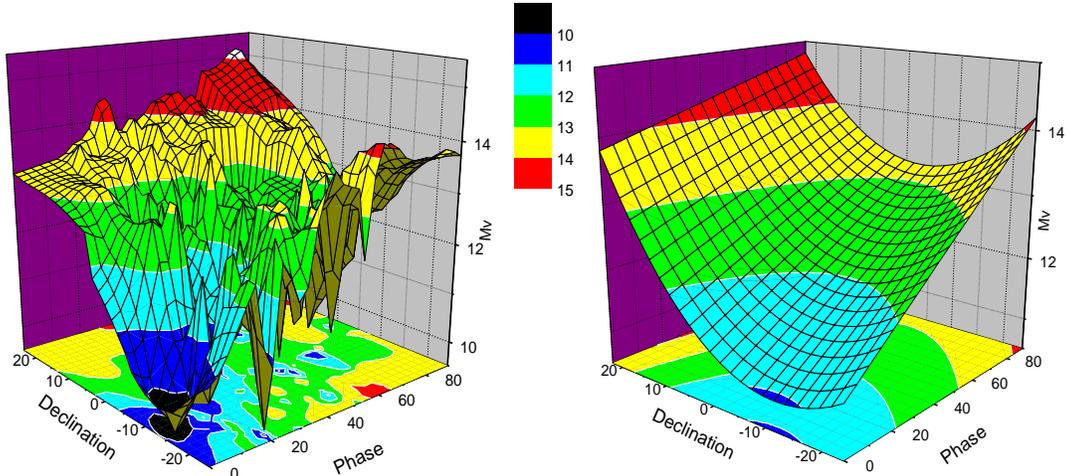


Figure 10: Fit results for unstabilized HS-376 case

The stabilized HS-376 case presents the most interesting results. Figure 11 shows the more typically-observed situation: low-phase brightness decreasing with increasing phase angle, yet a considerable variation with declination angle. Localized glints appear as “stalactites” in the krigged plots. Figure 12 represents a more unusual but not

unique situation of high-phase glint predominating the results. Some brightness increase at low phase is observed, but it is dwarfed by the considerably brighter response at high phase. This response is an artifact of the satellite construction, but it is consistently observed and thus merits modeling as such. Fit descriptive data, given in Figure 13, shows again the surprising deference phase angle pays to declination angle, both as a first- and second-degree term. This case also manifests the largest error statistics, requiring a full magnitude's augmentation of the fit's constant term to reach a 95% minimum brightness level.



Figures 11 and 12: Krigged and fit data for stabilized HS-376 objects 21906 (above) and 21964 (below)

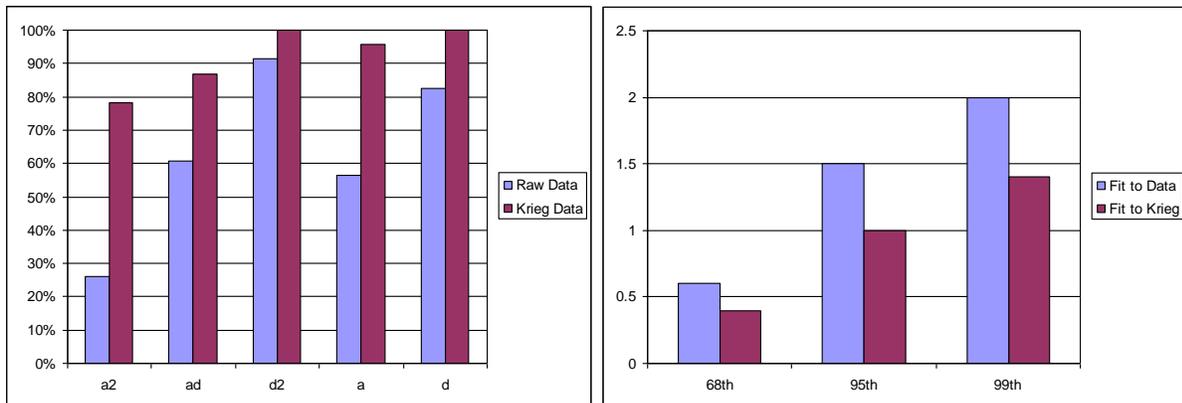
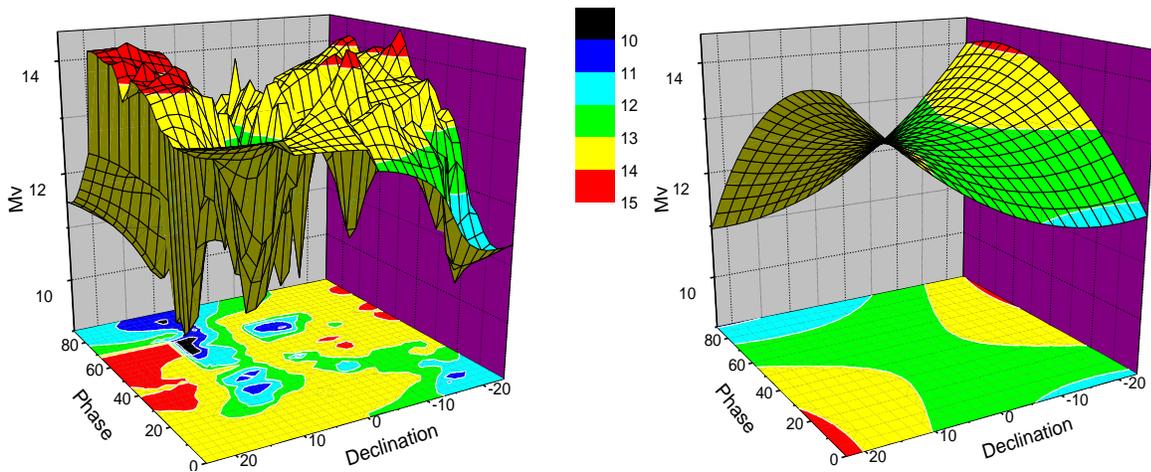


Figure 13: Fit descriptive data for stabilized HS-376

6. CONCLUSION AND FUTURE WORK

The optimal properties of kriging make it the gridding method of choice for visual examination of the functional relationships between observed satellite brightness and the viewing geometry, principally solar phase angle and solar declination angle. Examination of a trial dataset argued for the use of a second-degree fitting function, which performed well in terms both of visual matches to the krigged data and reasonable error statistics. The including of solar declination angle as an independent variable was seen as essential, as it affects the observed brightness more profoundly and more frequently than the conventionally-embraced solar phase angle. Fit behavior did not differ drastically between the data-sourced and krig-grid-sourced fits; but although space limitations for this paper did not permit presentation of some typical cases, it was observed as part of the analysis that the fits from krigged gridding data avoided extreme behavior at the edges of the fit-space and were thus the preferred solution. Quadric surface fits to the krigged data grid are thus the recommended procedure for future releases of the VM SATCAT.

One kriging data product that was not examined as part of this analysis is the set of error variances that can be calculated for each gridding point, primarily because the Origin kriging product did not export these data. However, if kriging is used as part of the VM SATCAT, these error variance data will be used at the least to produce for each satellite an error model to use in conjunction with the brightness model, allowing a regionalized rather than merely global error estimate. Another possibility is to maintain a database of krigged grid points and error variances, allowing direct interpolation of these data to predict brightness (and variance) for a particular viewing geometry.

Finally, some re-validation of the above will be necessary when plentiful DEEP STARE data are available. Ground-based sensors tend to produce more persistent and exaggerated glint responses, and the adequacy of the second degree model may need to be re-evaluated.

7. REFERENCES

1. Hejduk, M.D., Kervin, P.W., Lambert, J.V., Stansbery, E.G., Africano, J.L., and Pearce, E.C., "Visual Magnitude Satellite Catalogue Release 1.0", 2001 AMOS Technical Conference, Maui HI, SEP 2001
2. Lambert, J.V., "Interpretation of Geosynchronous Satellite Phase Angle versus Magnitude Relationships", Contract F05603-90-C-0010 Specialized Data Report: MOTIF FY95-01, 30 NOV 1994.
3. Lambert, J.V., "Analysis of Magnitude versus Phase Angle Data for Two Classes of Deep Space Satellites", Contract F05604-95-C-9011 Specialized Data Report: MSSS FY96-04, 11 OCT 1996.
4. Lambour, R.L., "Phase Angle Dependence of Satellite Brightness Derived from Space-Based Visible Data", MIT/LL Project Report SPC-8, 20 FEB 2001.
5. Lambour, R.L. and Sayer, R.W., "ETS Measurements of Satellite Phase Curves at High Phase Angles", MIT/LL Project Report ETS-138, 12 SEP 02.
6. Davis, J.C., *Statistics and Data Analysis in Geology*. New York: John Wiley & Sons, 1986.
7. Olea, R.A., Optimum mapping techniques using regionalized variable theory: *Kansas Geological Survey Series on Spatial Analysis*, no. 2. 1975.