SATELLITE BRIGHTNESS ESTIMATION USING KRIGING OPTIMIZED INTERPOLATION

Jennifer M. Okada
Boeing LTS Inc., 1250 Academy Park Loop, Colorado Springs, CO 80910
Matthew D. Hejduk
Titan Corporation, 10006 Willow Bend Drive, Woodway, TX 76712

ABSTRACT

With the advent of CCD optical technologies, the serendipitous collection of single-point photometry as part of metric tracking is now commonplace; and a number of applications for these data have been suggested and are in the early stages of experimental implementation. Many such applications, however, rely on the ability to predict the photometric brightness of a given satellite illumination geometry, whether to determine sensor probability of detection of a given satellite pass or in order to calculate a brightness "residual" between the observed and expected brightness values. Two-parameter brightness models are frequently employed in asteroid observation and, with some modification, have been applied to spacecraft, but these approaches relied entirely on deterministic modeling for an application that exhibits much stochastic behavior. The Kriging modeling approach, derived from mining and geological research, provides a unique way to combine deterministic and stochastic models and thus produce optimized (minimum-variance) gridding. This approach has been proposed and theoretically defended for spacecraft brightness estimation but has not been implemented in a manner appropriate for spacecraft and run against large photometry datasets.

The present project has performed and exercised such an implementation. The target dataset is first "detrended" using a straight-line deterministic model with glint accommodation. Next, phase-declination anisotropy is investigated and scaling factor(s) applied to the dataset, if necessary, to transform the dataspace into an isotropic situation. After this, a full-dataset variogram (the inverse of a correlogram) is computed from the isotropic dataset—this will determine the relationship between data point separation distance and the resulting amount of change in variance. Once the experimental variogram has been fitted to an appropriate model, the model results can be used to determine the degree to which any given observation can be said to be correlated with any other observation: this is the basis by which the brightness value at a certain grid point can be expressed as a weighted linear combination of the surrounding observations, with the weights to apply to each measurement determined from the variogram model. This approach produces both a brightness map over the desired grid space and additionally a variance map that gives a (relative) statement of the uncertainty of each gridded point. Several statistical tests are then run to assess the quality of the entire brightness-space solution for the object.

The particular implementation includes a GUI mode that allows the analyst to control the particulars of the detrending, anisotropy, and variogram solutions, as well as to file up the results of a satisfactory solution, yielding brightness gridding matrices that then can produce further brightness granularity with two-dimensional interpolation. It also includes a batch mode that, using system defaults, will run against an entire catalogue of brightness data for many objects. The latter is the functionality employed presently to produce brightness products for the new SP Sensor Tasker, operated by AFSPC/21stSW/1SPCS.

0. SATELLITE ILLUMINATION FUNDAMENTALS

The observed brightness of a satellite is a function of both the degree of illuminated surface that faces the observer and the reflective properties of the particular surfaces that happen to be illuminated. The accompanying diagram (Fig 1) illustrates the (angular) method by which these illumination conditions are usually described. The solar phase angle, which is the angle from sun to spacecraft to sensor, governs the amount of illuminated surface that will face the observer; the phases of the moon, for example, are examples of low (full moon), moderate (half moon), and large (new moon) phase angles. Once the amount of exposed illuminated surface is determined, two additional angles are necessary to define which particular parts of the satellite fall in the illuminated area facing the observer; while many angles could be chosen, two convenient angles are created by the definition of an aspect vector, with vertex at the satellite's center of mass and proceeding normal to the satellite's principal surface. The aspect angle is that formed by the aspect vector-spacecraft-sensor vector pair, and the obliquity angle is the dihedral angle between

the planes formed by the phase and aspect angles. With these three angles known, the illumination geometry can be specified completely.

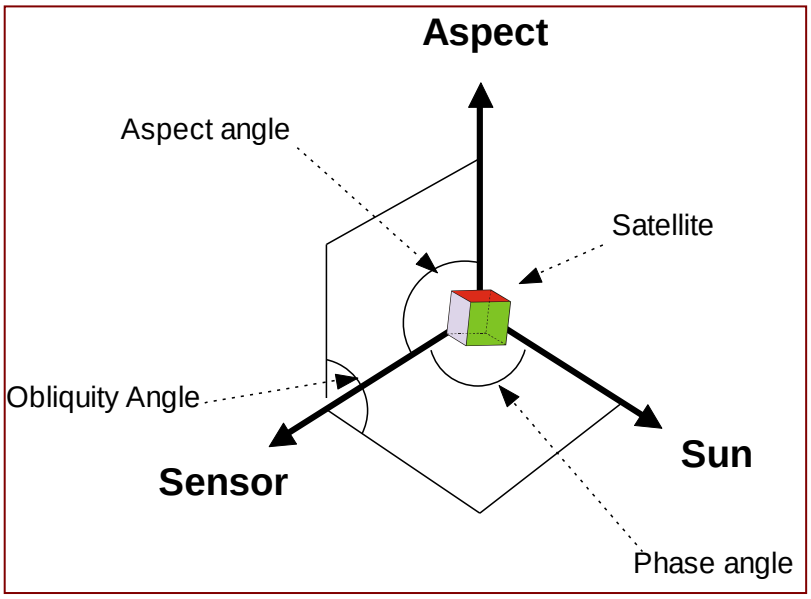

Fig. 1 Satellite Illumination Geometry

Unfortunately, for most spacecraft the particular satellite orientation is not known, making the construction of the aspect vector impossible and therefore leaving the aspect and obliquity angles indeterminate; and at first this would seem to counsel the use of phase as the only independent variable in attempting to predict satellite brightness. However, most geosynchronous satellites are three-axis stabilized and thus maintain a fixed orientation with respect to the earth (and ground-based sensors). In such cases, the aspect vector can be constructed; but one can proceed more easily in noting that the illuminated part of the satellite will be dictated entirely by the solar declination angle (determining whether the satellite is top- or bottom-lit) and the latitude of the ground-based sensor. In such cases, solar declination angle, which is very easily calculated as a function of the day of year, can replace the more complicated construction of the aspect vector and calculation of the resultant angles. In those additional cases in which there is a quasi-earth fixity (such as spin-stabilized geosynchronous satellites or certain other standard orbits), the use of the solar declination angle could be expected to help somewhat, so it was seen as expedient to advance both solar phase angle and solar declination angle as independent variables for the purposes of satellite brightness modeling.

I. SOFTWARE INITIAL SETUP

In GUI mode, the Kriging Optimized Interpolation (KOI) Tool software initially displays the following setup window (Fig. 2). In this window, the user can select the satellite of interest, observation dates, and the photometry database to be queried (data are presently available from the three GEODSS sites, the RAVEN telescope at MSSS, and SBV). Photometry data can also be loaded from a stand-alone ASCII file. The "Start" button moves the program to the next process, which is the application of a deterministic model to "detrend" the data. The user also has the option to skip all the intermediate windows and display only the final results window.

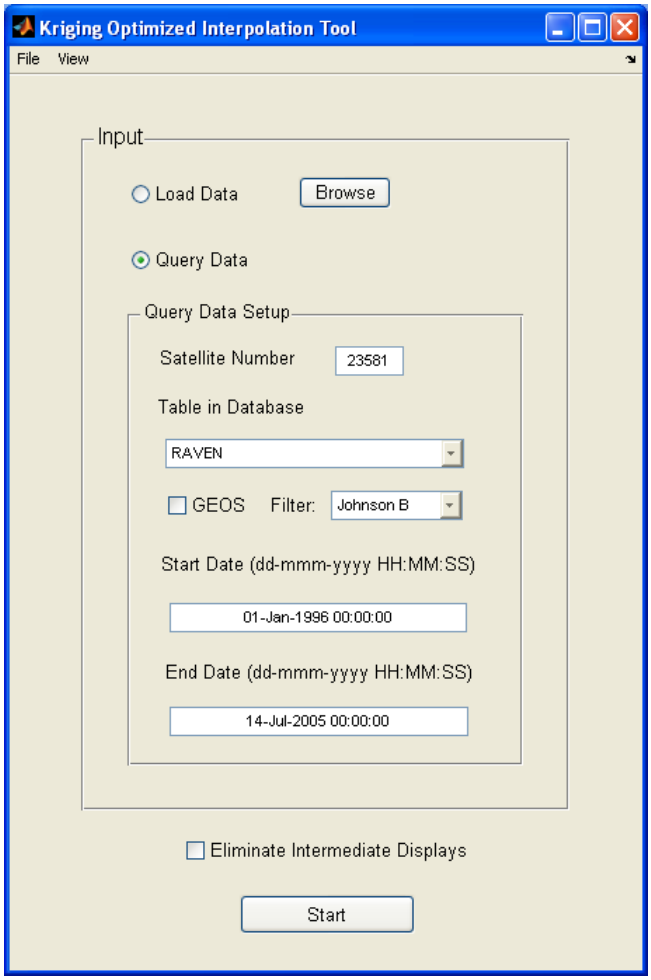

Fig. 2 Setup Display

II. DETRENDING THE DATA

KOI produces brightness predictions through the application of two distinct model types, a

deterministic model and a stochastic model. Since it is rare to be able to produce a high-fidelity satellite deterministic model that will yield reliable brightness predictions, the present strategy is to apply a generic deterministic model to the data and try to fit the resultant residuals to a stochastic model, thus hoping to account for the entire photometric response through the simultaneous use of both models. The application of the deterministic model in order to remove systematic effects, leaving only residuals to be modeled stochastically, is called "detrending." A number of previous studies have established, with reasonable error boundaries, a linear relationship between satellite brightness and solar phase angle [1, 2. 3, 4, 5]; so this seems a natural model to apply, especially in the absence of any obvious deterministic relationship between brightness and solar declination angle (since such a relationship would depend entirely on the construction details of a particular spacecraft). The data are divided into a positive and negative phase angle group, further divided into 3-degree phase bins, and each bin then summarized by the 68th percentile value of the brightness values in that bin. A robust regression fit is then applied to the binned data summary points, and a t-test to the resultant slope is conducted to determine if the resultant linear behavior is statistically significant. If it is, then this resultant line is used as the deterministic model; if it is not, then no deterministic model is applied; and for this satellite one relies entirely on the stochastic model. The fit results are displayed in the detrend window (Fig. 3). If the fit is successful, the deterministic model values are subtracted from the dataset, leaving a residual set to be accommodated by the stochastic model as part of the Kriging step. If one half-plane produces a statistically-significant fit and the other does not, the mean of the data is subtracted from the non-fit side in order to create compatibly-sized residuals across the entire phase-declination space.

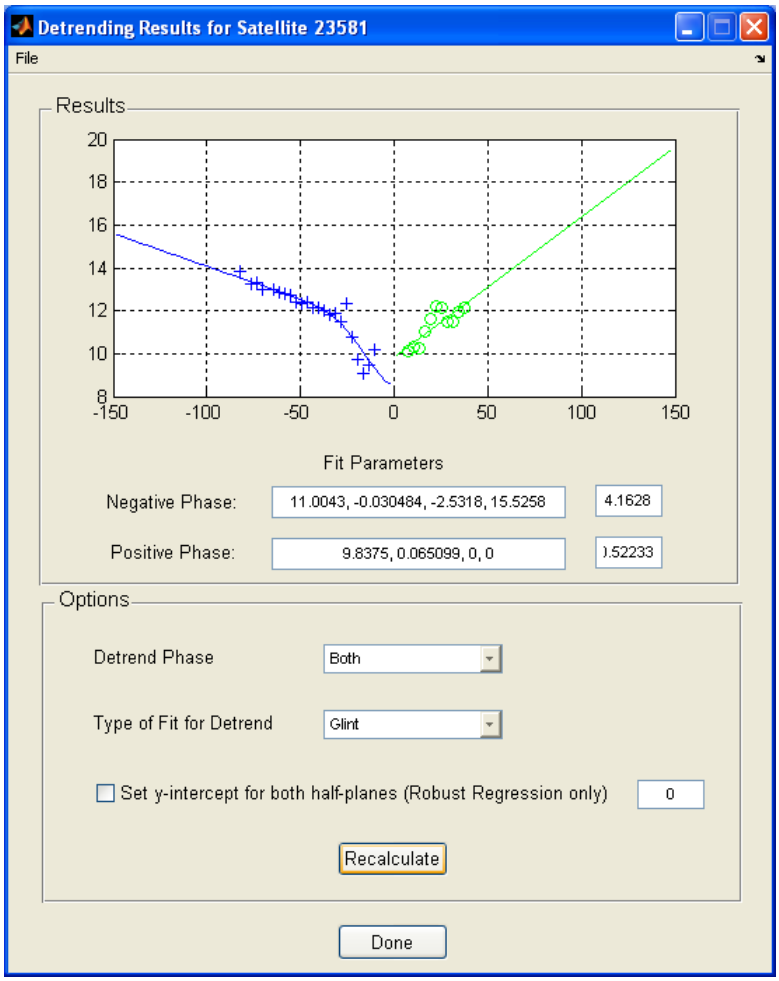

Fig. 3 Detrend Display

While satellite photometric returns are almost always a combination of diffuse and specular response, the specular component can become very large and predominate the response at low phase angles [3]. Since this is an anticipated effect, an attempt should be made to accommodate it by the deterministic model. The software tests for the presence of glint response by comparing the individual bin summary brightness values (for phase angles less than 20 degrees) to that predicted by the robust regression line; and if any of these differences exceeds one visual magnitude, the presence of glint response is presumed. In glint situations, the bin summary values are fitted to a model that is a combination of a straight line and "Gaussian" response, of the form

$$M_v = A + B\alpha + C\exp\left(-0.5*\left(\frac{\alpha}{D}\right)^2\right)$$

in which α is the phase angle and A-D are constants to be determined by the fit. In principle, this form of the model could be used in all cases, with constant C assuming a value near 0 for cases that exhibit no glint. However, the uncertainties and instabilities of non-linear optimization counsel, at least in batch mode, the invocation of this more elaborate model only when it appears to be warranted.

The GUI allows the user to enable or disable the glint model, pick the method of fitting, control whether or which half-plane is actively fit, and force both half-plane fits to share a specified common y-intercept. Selecting "Done" moves to the next analytical phase of the program, which is the isotropic transformation.

III. TRANSFORMING TO ISOTROPIC CONDITIONS

As was discussed in the initial section, the stochastic brightness model is constructed to be a function of both solar phase angle and solar declination angle. One can imagine (and indeed the result will take this form) a grid with phase on the *x*-axis and declination on the *y*-axis, and observed (and predicted) brightness as the out-of-plane (*z*) variable. Constructing an algorithm for an optimized interpolation of this grid is straightforward if the situation is isotropic, that is, the difference in brightness between two points is independent of the direction of the segment between them. In such a situation, the expected correlation between the brightness of two points would be the same both purely in the phase direction and purely in the declination direction.

Of course, there is no reason to suppose, and good *a priori* (and empirical) reasons to doubt, that such a situation inheres for most satellites. The hope, however, is that through a linear transformation (rotation and scaling of axes) this anisotropic situation can be transformed into an isotropic one. Since in the present case there is no ambiguity from model definition regarding the orientation of the axes (as there might be, for example, if *x* and *y* were spatial variables whose alignment with physical phenomena needed to be secured), the only variable to determine is a scaling factor (here chosen to be applied to the *y* or declination axis) To calculate this anisotropy scaling factor, one must determine the degree of brightness correlation in each discrete direction and then the scaling factor that will minimize the inter-directional difference of this correlation.

The usual method in regionalized variable theory for measuring correlation is actually to measure its inverse, which is the degree of variation as a function of separation distance, the so-called *variogram*. The equation for the *raw* variogram is given in Eq. 1. It is a scatter plot of the square of brightness (*z*) variation versus distance (*h*) between data points:

$$\frac{1}{2}\left[z(\vec{x}_i) - z(\vec{x}_i^{\,/})\right]^2 \; vs\; h, \quad h = \left\|\vec{x}_i - \vec{x}_i^{\,/}\right\|$$

Eq. 1

The *experimental* variogram, given in Eq. 2, bins the data points and gives a functional relationship:

$$\hat{\gamma}(h_k) = \frac{1}{2N_k}\sum_{i=1}^{N_k}\left[z(\vec{x}_i) - z(\vec{x}_i^{\,/})\right]^2, \quad h_k^l \le \left\|\vec{x}_i - \vec{x}_i^{\,/}\right\| < h_k^u$$

Eq. 2

An example of an experimental variogram plot, specifically one that constitutes what is called a stationary function, is shown in Fig. 4. The semivariance (γ), which gives a measure of degree of variation, approaches zero at small separation distances (h) and increases with separation distance, until finally it approaches a stationary value, or sill. At this separation distance (a, the correlation length), the correlation between points becomes the same as the overall variance of the sample and therefore no longer increases with increased separation distance; at this distance, brightness values are no longer considered correlated and the boundary of the region of the "regionalized variable" has been reached. Not all variograms have the same appearance as Fig. 4; some never reach a sill (a non-stationary situation), and others maintain some intrinsic variability even at zero separation distance (the so-called "nugget effect"); but experience shows that a variogram with the above characteristics is a reasonable expectation for satellite brightness data.

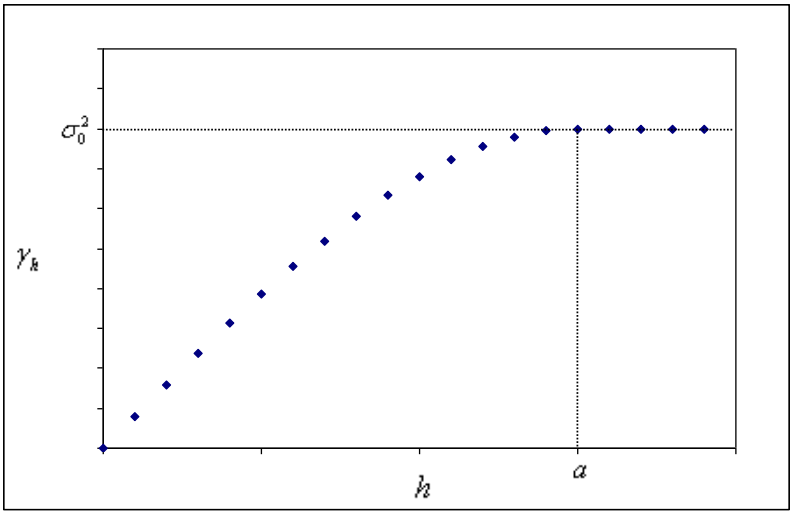

Fig. 4 Experimental Variogram

To determine the variation in correlation as a function of direction (*i.e.*, the anisotropy), we construct variograms in each direction and compare them. This is done by dividing the phase-declination plane into "wedge angles" with their vertices at the origin and considering only those points that fall within each wedge angle for the calculation of the variogram in that direction. Choosing a wedge angle of thirty degrees results in twelve wedges, each with its own variogram (presuming each wedge contains enough data to permit variogram calculation). The feature of interest from each variogram is the correlation length a, the length at which the sill is reached.

The ordered pairs of wedge angle and correlation length are plotted in polar coordinates. If the data were isotropic, the resulting plot would be a circle because all the correlation lengths would be equal. For anisotropic data, we attempt to recreate this situation by fitting these points to an ellipse using ordinary least squares and then determining the scaling of the declination axis necessary to transform the ellipse into a circle [6]. The software as presently implemented tries four different wedge angles of 15, 30, 45, and 60 degrees, using the wedge angle

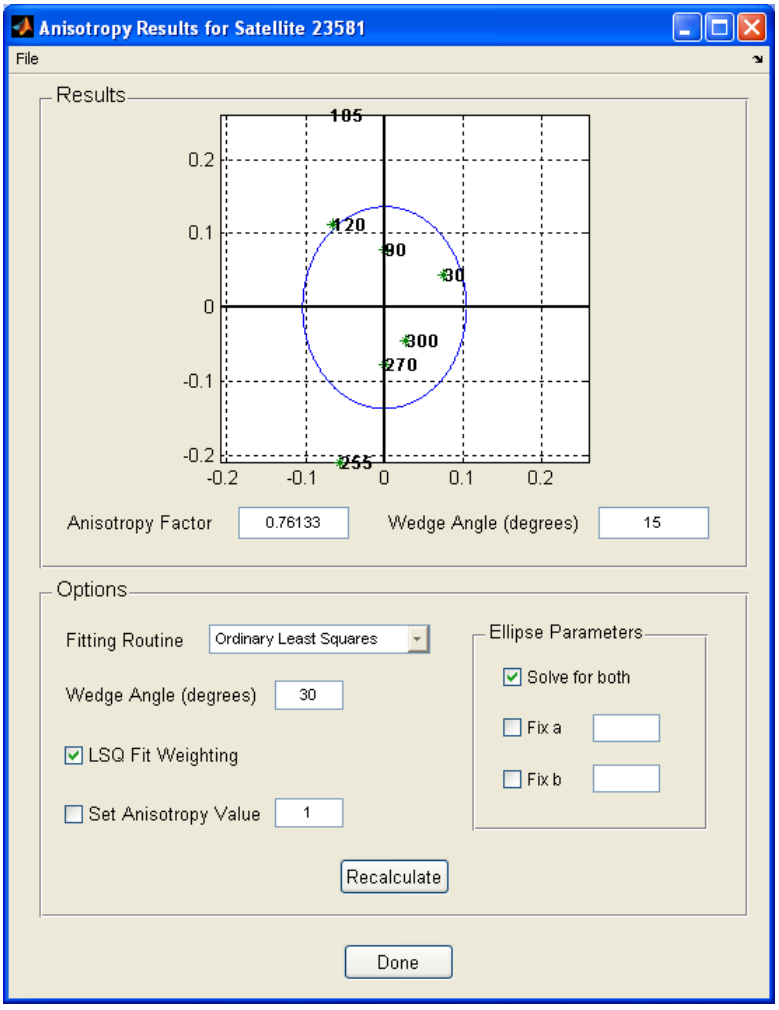

Fig. 5 Anisotropy Display

that produces the smallest residual sum of squares for the anisotropy factor calculation. Fig. 5 shows the results display for this process.

The user can choose to recalculate the anisotropy factor with different options. Instead of using an ordinary least squares fit, robust regression and weighted least squares are also available. The user can also set the wedge angle and one or both of the ellipse parameters, or skip the whole process and set the anisotropy factor himself/herself. The anisotropy factor is then applied to the dataset, transforming it into this new (mostly isotropic) coordinate system.

IV. VARIOGRAM CALCULATION AND MODEL FITTING

Now that an isotropic (or near-isotropic) situation has been produced via a coordinate transformation, a variogram can be calculated for the overall dataset. This variogram will provide the correlation-vs-separation-distance relationship needed for the optimized interpolation. As each desired brightness grid point can be presumed to be determined by a linear combination of the surrounding points, the variogram characteristics will determine how to weight those surrounding points in the linear combination. To use the variogram in a calculation context, we must curve-fit a canonical model to it. Models commonly used in variogram analysis and included in this software implementation are the Spherical, Exponential, Gaussian, Linear, and Power models, as defined by the equations below:

$$\begin{aligned}
&\textit{Spherical} && \gamma_h = C_0\left[\frac{3}{2}\frac{h}{L} - \frac{1}{2}\left(\frac{h}{L}\right)^3\right] + \gamma_0 \\
&\textit{Exponential} && \gamma_h = C_0\left[1 - \exp\left(-h/L\right)\right] + \gamma_0 \\
&\textit{Gaussian} && \gamma_h = C_0\left[1 - \exp\left(-\left(h/L\right)^2\right)\right] + \gamma_0 \\
&\textit{Linear} && \gamma_h = C_0 h + \gamma_0 \\
&\textit{Power} && \gamma_h = C_0\left(h/L\right)^{\lambda} + \gamma_0
\end{aligned}$$

Eqs. 3-7

The constants for each of the models above are determined by non-linear least squares fitting to the experimental variogram, and the model producing the smallest residual RSS is selected by the software. Results are shown in the Variogram Model Results window (Fig. 6.).

The user has the option to choose a different model and/or set the various model parameters. By clicking the "Update" button, the user can view the subsequent changes to the model. Clicking "Done" invokes the Kriging stage of the program.

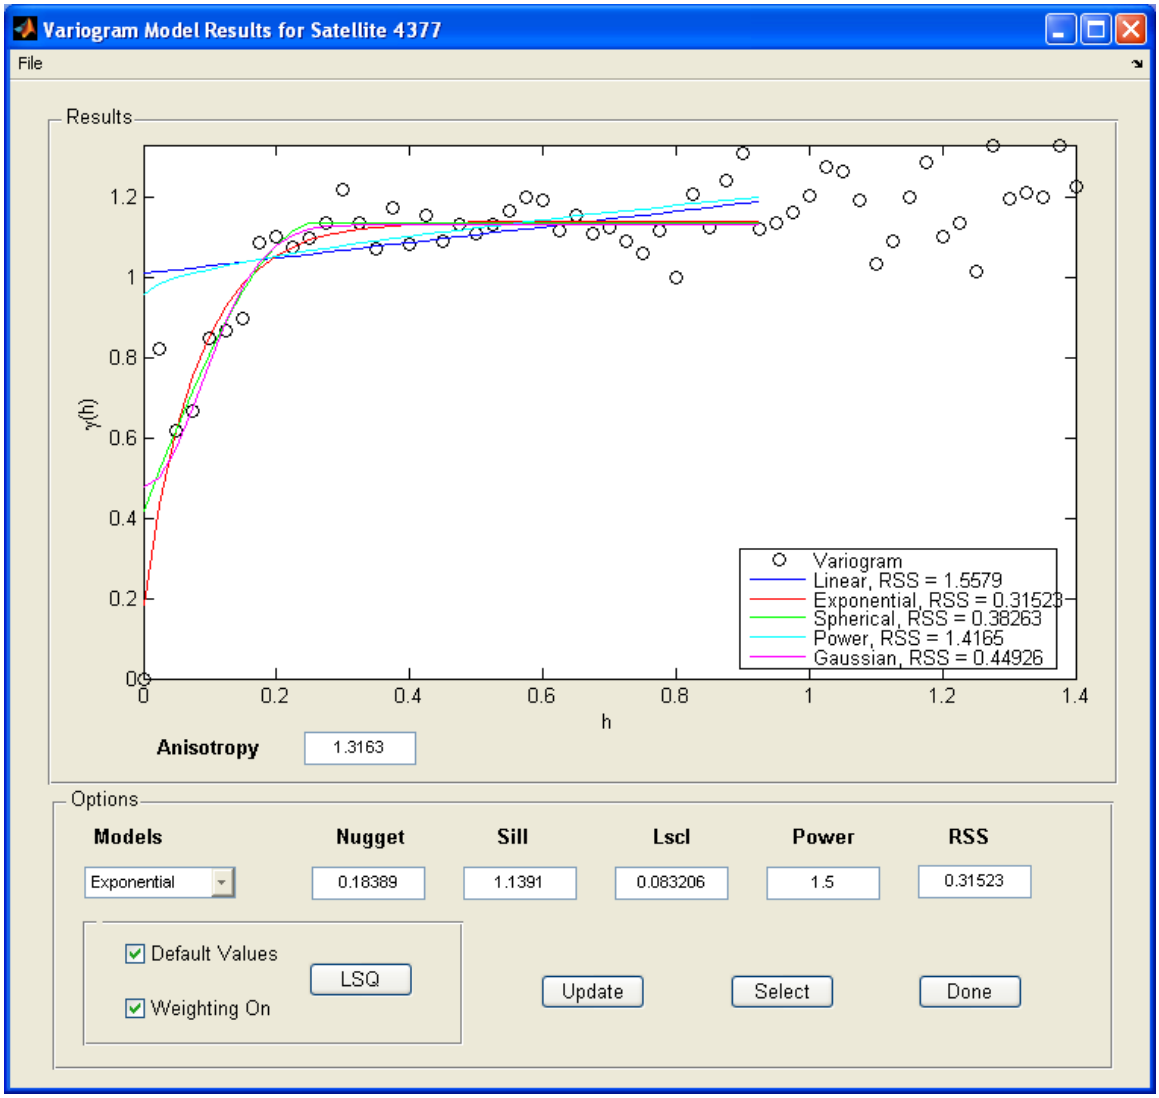

Fig. 6 Variogram Model Display.

V. KRIGING

Kriging estimates the brightness at each point in the brightness grid by presuming this brightness to be a linear combination of the surrounding observed data points. Each point used in the linear combination is weighted according to the variogram model, which was selected in the previous step. The constraints used in determining the set of weighting factors for the estimation are that the solution be unbiased (Eq. 8) and the variance be minimized (Eq. 9) [7]. Lambda represents the to-be-calculated weighting factors to be applied to the surrounding points, and gamma is the variogram function:

$$\sum_{i=1}^{n} \lambda_i = 1$$

Eq. 8

$$E\left[\left(\hat{z}_0 - z(\vec{x}_0)\right)^2\right] = -\sum_{i=1}^{n}\sum_{j=1}^{n}\lambda_i\lambda_j\gamma\left(\left\|\vec{x}_i - \vec{x}_j\right\|\right) + 2\sum_{i=1}^{n}\lambda_i\gamma\left(\left\|\vec{x}_i - \vec{x}_0\right\|\right)$$

Eq. 9

The software uses a Lagrange multiplier technique to effect the minimization solution.

A desirable additional feature of Kriging is the ability to generate an estimation variance unique to each determined gridding point, as a function of the variogram, the calculated weighting factors, and the Lagrange multiplier (v):

$$\sigma_0^2 = E\left[\left(\hat{z}_0 - z(\vec{x}_0)\right)^2\right] = -v + \sum_{i=1}^{n}\lambda_i\gamma\left(\left\|\vec{x}_i - \vec{x}_0\right\|\right)$$

Eq. 10

VI. MODEL VALIDATION

Two approaches to model validation are accommodated by the developed software: a visual approach, which uses a jackknife plot, and explicit statistical tests for model adequacy. The jackknife model validation sequentially removes one point from the data and predicts the brightness for that point from the remaining data using the Kriging technique. The jackknife display (Fig. 7) plots each observed point (red x's) and about each a high and low predicted point, derived from the prediction plus and minus a factor of the calculated Kriging variance (blue x's). Plots sorted on phase and declination, as well as a purely sequential plot, are provided. The display is somewhat difficult to read at less than full size, but it allows the user visually to determine how often the observed value falls within a given range about the predicted value.

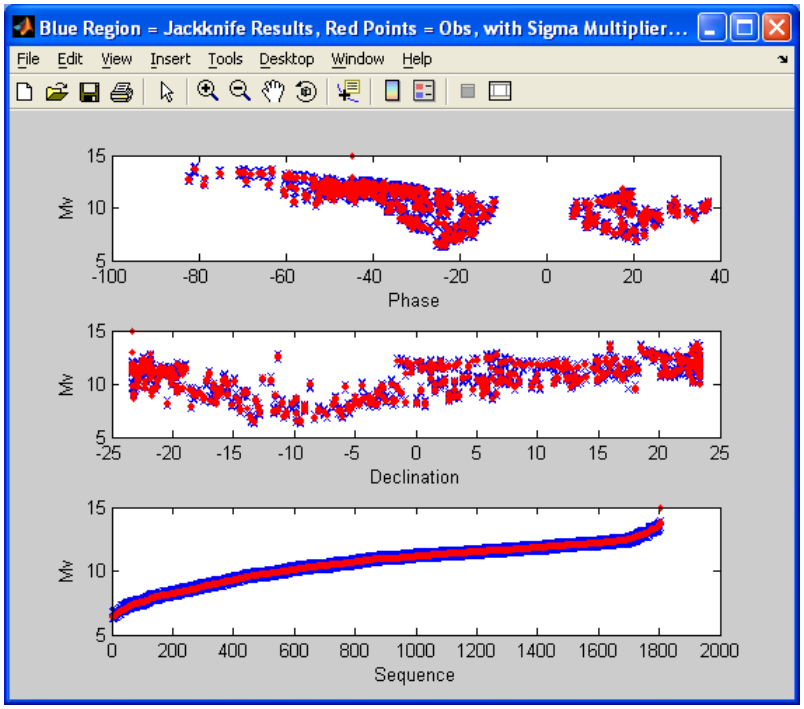

Fig. 7 Jackknife Display.

In addition to visual assessment tools, a number of statistical tests are run to help determine the adequacy of the constructed model. The two attributes we wish primarily to test are unbiasedness and adequacy of error modeling. To execute such tests, we first need to construct a set of residuals to which to apply them. A convenient way to do this is to order the data points at random and, beginning with the first data point, predict the value of the second and calculate the (residual) difference; then use the first two to predict the third and calculate a residual, &c. until the entire dataset has been so analyzed and a residual set produced [7]. The so-called "Q_1" test assesses whether the residuals are biased, and it does this by taking the mean of the residuals, as shown in Eq. 11 below:

$$Q_1 = \frac{1}{n-1}\sum_{k=2}^{n}\varepsilon_k$$

Eq. 11

Presuming all the systematic errors have been solved for (an assumption to be discussed explicitly below), one can thus expect a Gaussian distribution with 0 mean and variance 1/(n-1), where n is the number of residuals. From this information, the p-value for Q_1 can be calculated and therefore the reasonableness of the unbiasedness presumption assessed.

If the calculated Kriging variance models the error well, the ratio of the calculated residual to the square root of the Kriging variance should be 1. A convenient way to test for this is to presume that this ratio is a normal variable; if so, then the sum of the squares of this ratio should conform to a chi-squared distribution with n-1 degrees of freedom, constituting the "Q_2" test below:

$$Q_2 = \frac{1}{n-1}\sum_{k=2}^{n} \varepsilon_k^2$$

Eq. 12

As with the Q_1 test, p-values for the Q_2 factor can be calculated and the degree of conformity to the chi-squared distribution determined.

Finally, and especially since the Q_1 and Q_2 tests implicitly require it, these residuals themselves should be tested for conformity to normality. There are many appropriate goodness-of-fit tests described in the literature, but two that have proven both robust and convenient from an implementation perspective are the D'Agostino D-test (variant of Shapiro-Wilk) and an omnibust 3rd and 4th moment test that examines the departure of the skewness and kurtosis from that expected from a normally-distributed population [8]. The results from all four of these tests are displayed as part of the results window, described below.

VII. RESULTS DISPLAY

The results window (Fig. 8) displays all the results from the previous windows along with the brightness and variance maps, vital satellite statistics, and a data density plot. The model validation results for the statistic and normality tests are also given. The jackknife display window can be opened by clicking the "Jackknife" button. By selecting "Save" from the "File" menu, a user can save the brightness and variance maps to a text file.

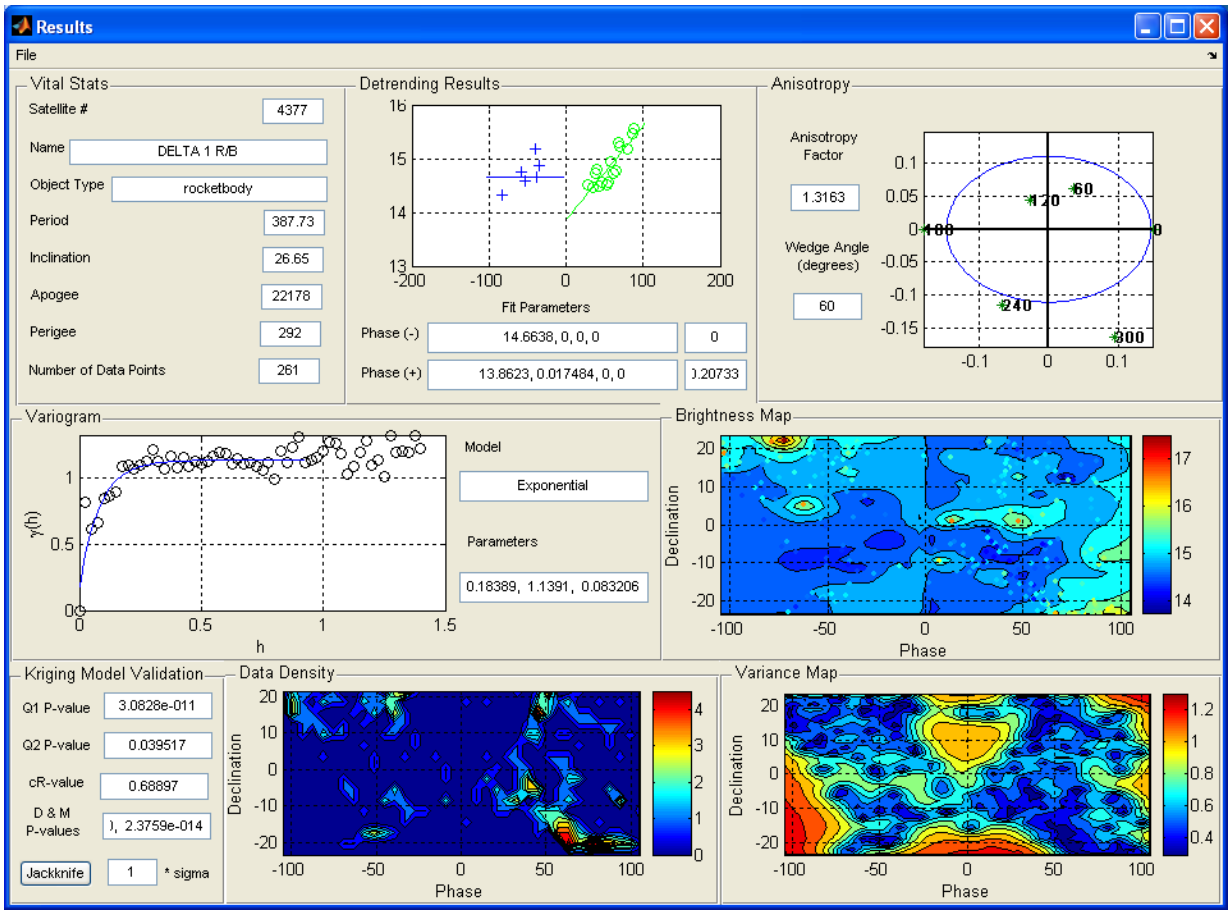

Fig. 8 Results Display

VIII. BATCH MODE

A batch run can be implemented by selecting "Batch Window" from the "View" menu on the initial window. This opens a window (Fig. 9) where the user can setup a batch run. The user can select a list of satellites to process by entering or browsing for the file. The user can also select the file containing a list of wedge angles to use for the anisotropy calculation and whether or not to execute a full run for all the wedge angles or only the best angle. The user can select the database and minimum number of observations for a satellite to be processed. The user can also decide whether or not to detrend the data and, if so, check for glints.

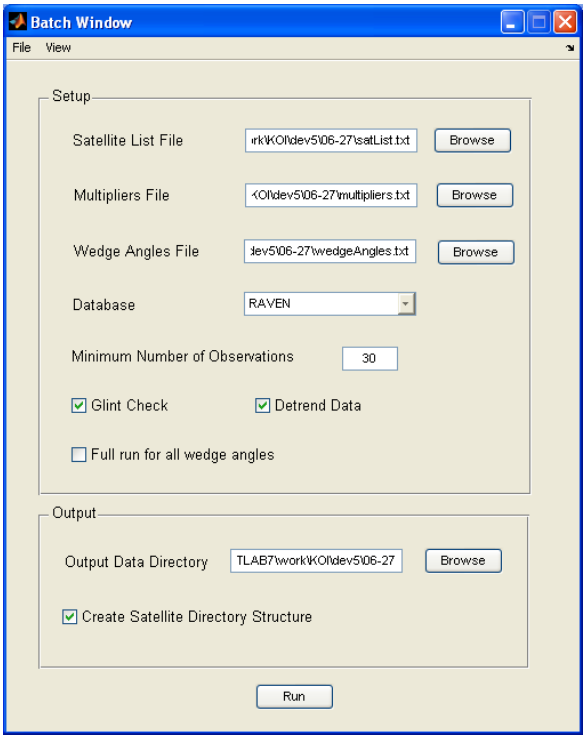

Fig. 9 Batch Mode Display

IX. FUTURE WORK

We are currently working on functionality to apply the Kriging optimized interpolation process to color photometry data. The present implementation already includes a feature to produce brightness and variance maps for color photometry within a single color band. The extension of the present functionality would be to use detrending and

variogram parameters from a satellite's "white-light" solution, which is generally a more data-rich environment, to perform interpolation for relatively data-poor color bands. It is not clear at present whether it is reasonable to presume these parameters to be relatively invariant across a satellite's color bands, but a fairly straightforward study effort should determine this.

X. 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. Journel, A.G. and Huijbregts, Ch. J. *Mining Geostatistics*. London: Academic Press, 1978.

7. Kitanidis, P.K. *Introduction to Geostatistics: Applications to Hydrogeology*. Cambridge: Press Syndicate of the University of Cambridge, 1997.

8. D'Agostino, R. and Stephens, M. Goodness-of-Fit Techniques. New York: Marcel Dekker, Inc., 1986.