

Kinematic-Vision Residuals Analysis

Kevin Nickels
Trinity University
Department of Engineering Science
One Trinity Place
San Antonio TX 78259-7200
210-999-7543
knickels@trinity.edu

Max Bajracharya, Ashitey Trebi-Ollennu, Robert Liebersbach
Jet Propulsion Laboratory
California Institute of Technology
4800 Oak Grove Dr.
Pasadena, CA 91109, USA
maxb@robotics.jpl.nasa.gov, ashitey@jpl.nasa.gov,
Robert.R.Liebersbach-121188@jpl.nasa.gov

Abstract—This paper describes the design, validation, and integration of a tool to locate a portion of the Instrument Deployment Device (IDD) on the Mars Exploration Rover (MER) vehicles on Mars in imagery from the front Hazard avoidance cameras, and to track the differences between the predicted and detected position of the manipulator over time.

The analysis of Kinematic-Vision Residuals, or the difference between where a manipulator is expected to appear in on-board imagery and where it actually appears in the imagery, yields insight into several aspects of an operational robotic system. The fidelity of the IDD and camera models is evaluated. Systematic changes in the performance over time can give insight to rover degradation or other changes. Finally, new models can be proposed and evaluated on the basis of trended data over time.

TABLE OF CONTENTS

- 1 INTRODUCTION
- 2 THE FIDUCIAL DETECTOR
- 3 ARM & CAMERA MODELS
- 4 RESULTS
- 5 CONCLUSIONS
- 6 FUTURE WORK

1. INTRODUCTION

The Mars Exploration Rover (MER) vehicles were designed to deploy scientific instruments on Mars [1]. To support this mission, several engineering cameras are located on the rover, including the front hazard-avoidance cameras (haz-

cams). The hazcams are utilized to avoid hazards while driving autonomously, as well as to select science targets to sample with one of the three in-situ instruments located on five degree-of-freedom robot arm (known as the Instrument Deployment Device, or IDD) that is used to place three in-situ instruments: the Alpha Proton X-Ray Spectrometer, or APXS, the Mössbauer spectrometer, and a microscopic imager, or MI.[2].

Since science targets are selected in imagery from the hazcams, the difference between the 3D location of a target selected in this imagery and the actual 3D location reached by the IDD is a useful metric. This difference, termed the *Kinematic-Vision Residual*, is formally defined to be, for a given pose of the IDD, the root-sum-square (RSS) difference between the 3D location of a point on the IDD predicted by the IDD model and the 3D location of that same point as measured by the hazcams and their associated models. Neither part of this closed loop is perfect - there are inevitable errors in both the IDD model and the camera models, and this analysis does not subscribe errors to either the IDD or the camera system but rather looks at the closed loop system performance.

2. THE FIDUCIAL DETECTOR

The Mössbauer spectrometer has a contact plate (MBCP) that is visible in some images, such as Figure 1, that are taken when the rover is positioning and utilizing the MI. Although not designed to be a fiducial, because the MBCP has inner and outer rings with well known dimensions and tend to appear in hazcam images with contrasting intensities, an algorithm to detect the plate can be designed.

To achieve maximum robustness and accuracy, the approach taken was to match a directed edge template of the fiducial by




Figure 1. An example pose from the Spirit MER vehicle with the MBCP indicated.

shifting a template of the MBCP with increasing step resolution in a local area and of the image. The template is created by projecting the edges of the MBCP rings to the images using the kinematic prediction of the 3D position of the MBCP in space.

Although using a feature detector (a Harris corner detector [3] for example) is fast and invariant to affine transformations, it is prone to false positives, particularly with direct sunlight on the specular metal end-effector. Some examples of images that challenge fiducial detection are given in 2.

With a good prediction, a standard image-based template correlation method (using normalized cross correlation, [4]) would be robust to false matches like spurious points or edges, but is computationally expensive. A further motivation for using the model-based correlation approach is that, while fiducials were not explicitly placed on the MER IDD, portions of the instruments on the end-effector can serve the purpose but are difficult to detect with a feature detector. The contact ring of the Mössbauer Spectrometer appears as the projection of two concentric circles and provides a unique target with respect to the environment and spacecraft visible by the cameras. However, unlike a standard corner fiducial it has the disadvantage of requiring foreshortening compensation to accurately localize its center.

The algorithm implemented is a correlation-based approach, but only matches step features (edges) to improve robustness and speed. Using the camera, IDD, and fiducial models, a template of sampled fiducial edges is created and then cor-




Figure 2. Some example poses illustrating the variability of appearances of the MBCP. Note the specular reflections in some images.

related across the image. Although of the same computation order as correlating an image-based template, correlating sampled edge points produces a significantly lower multiplier constant.

For a given IDD pose, the fiducial template is created by sampling the predicted position of the rings' edges (circles of radius r_{inner} and r_{outer} , as angle θ ranges from 0 to 2π) and projecting these positions to 2D image coordinates (p_s) via camera and IDD models (M_{world}^{camera} and $T_{fiducial}^{world}$). More information about the camera and IDD models can be found in [5] and [6], respectively.

$$p_s = M_{world}^{camera} T_{fiducial}^{world} \begin{bmatrix} r \cos \theta \\ r \sin \theta \\ 0 \end{bmatrix} \quad \forall \theta \in [0, 2\pi], r \in \{r_{inner}, r_{outer}\} \quad (1)$$

The camera model accounts for all image parameters (including distortion) and consequently works for very wide field of view images without requiring the expense of image rectification. While the parametric template implemented is specific to the MER fiducial, any template with edges or step functions could be used.

The correlation score is defined as the sum of directional derivatives (tangent to the edge) at each of the image (I) coordinates in the direction of the fiducial center (c).

$$c = M_{world}^{camera} T_{fiducial}^{world} \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix} \quad (2)$$

To find the fiducial, the template is correlated across a local area (of window size w) in the image with shifts of size Δp . Initially, the template is correlated at only the pixel level ($\Delta p = 1$). Once the maximum score has been found (at $\Delta \bar{p}$) the template is correlated at sub-pixel steps in a sub-window about this point to localize the template further.

$$\Delta \bar{p} = \max_{\Delta p} \sum_{s+\Delta p} \left[I\left(p_s - \frac{p_s - c}{\|p_s - c\|}\right) - I\left(p_s + \frac{p_s - c}{\|p_s - c\|}\right) \right] \quad \forall \quad \Delta p \in [-\lceil w/2 \rceil, \lceil w/2 \rceil] \quad (3)$$

An example of the initial sub-sampled template points and the locations that maximize the location are shown in Figure 3 b. The correlation score using $\Delta p = 1$ and $\Delta p = 0.1$ are shown in Figure 3 c and d. Note that the curvature of the peak of the sub-pixel step size is an indication of the precision that can be achieved.

Verification of Fiducial Detection Results

This fiducial detector, as designed, is quite reliable. It can find the MBCP in a variety of situations, as illustrated in Figure 4. However, the false-positive rate of the fiducial detector as initially implemented was too high for deployment as an automatic tool in the downlink processing for the MER vehicles.

To achieve even more reliable results, several metrics were considered. These include the intersection of the normal of the MBCP with the image plane, the contrast between the inner and outer rings of the MBCP, the aspect ratio of the MBCP in the images, and the *ray gap* (distance of closest intersection of the projection rays from the left and right cameras).

A random sampling of full-resolution fhazcam images was used to train the fiducial detector. Several observers independently rated the fiducial detector, recording both type I and type II errors (false positives and false negatives), as shown




Figure 4. Some example poses illustrating the performance of the fiducial detector.

in Table 1. A combination of the above metrics were used to reduce the type I error rate to less than 5%, one of the acceptance criteria. The type II error rate was not considered in the final testing, but every attempt was made to minimize it without unduly compromising the type I rate. For each data set, the first data column shows the number of poses in the data set. The second and third data columns show the number of poses where the fiducial detector correctly found the fiducial, as indicated by automatic and manual measures. The fourth through seventh columns indicate the number of true positives (the automatic system and manual scores both indicate that the fiducial detector succeeded in finding the fiducial), true negatives (the automatic system and manual scores both indicate that the fiducial detector did not succeed in finding the fiducial), false positives (the automatic system indicated that the fiducial detector succeeded but the manual scores indicated that it did not) and false negatives (the automatic system indicated that the fiducial detector did not succeed but the manual scores indicated that it did). Finally, the type I error rate is given for each data set. Note that the data set includes all possible images where the MBCP is viewable by the cameras, but does not account for images that are saturated, contain specularities, or have extreme viewing angles. This results in the low number of fiducials found as compared to the total number of poses.

An independent random sample of images was used to test the fiducial detector for acceptance. The type I errors were manually scored as described above. The results, as shown in Table 2, have type I error rates of under 5% for all data sets.

3. ARM & CAMERA MODELS

When projecting from joint angles to a 3D position in space, a deflected-kinematic model [7] that accounts for gravity-induced deflections as well as the calibrated forward kinematics of the IDD. When triangulating from the projection of the MBCP in the left and right image to compute the workspace location of the MBCP, a CAHVORE camera model [5] is

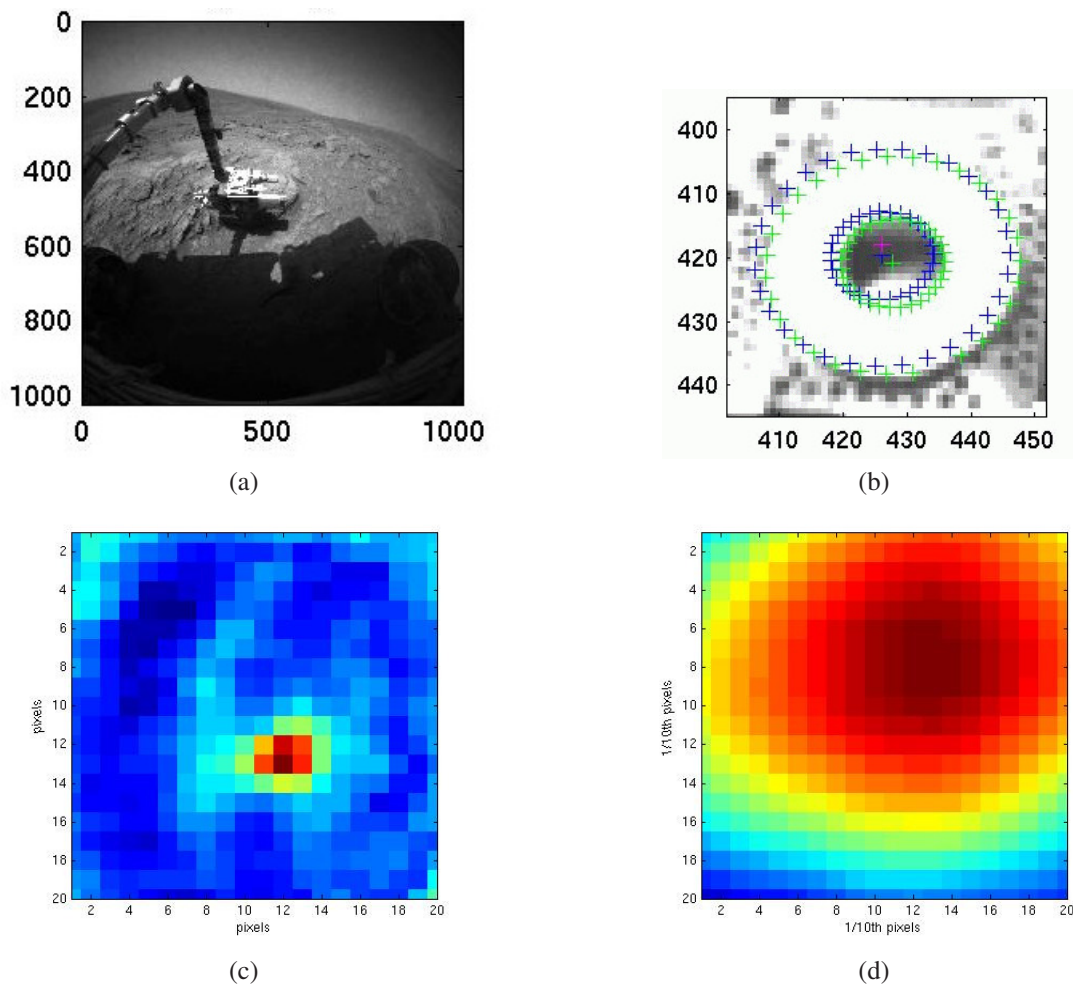


Figure 3. Pixel (c) and sub-pixel (d) correlation scores for the predicted and detected sampled template (+’s in b) of image (a). Red is a low correlation score (good match), and blue is a high correlation score (poor match).

Table 1. Training Data for Fiducial Detection

Rover	DataSet	Number Poses	Fiducial Found (Automatic)	Fiducial Found (Manual)	True Positives	True Negatives	False Positives	False Neg Negatives	1-FP Rate
mera	alpha	20	11	11	10	8	1	1	95%
merb	alpha	20	15	15	14	4	1	1	95%
mera	beta	17	4	11	3	5	1	8	94%
merb	beta	40	27	32	24	5	3	8	93%
mera	gamma	34	26	33	26	1	0	7	100%
merb	gamma	43	6	18	3	22	3	15	93%

Table 2. Testing Data for Fiducial Detection. A sol is a Martian day.

		Poses	MBCP Found (Auto)	MBCP Found (Manual)	False Positives	False Positive Rate
mera	sol 0-160 (512) - hips	88	29	27	2	2.3%
merb	sol 0-160 (512) - hips	147	64	62	2	1.4%
mera	sol 0-160 (512) - onboard	88	14	13	1	1.1%
merb	sol 0-160 (512) - onboard	147	64	63	1	0.7%
mera	sol 176 (training) - hips	20	7	7	0	0.0%
merb	sol 155 (training) - hips	20	14	14	0	0.0%
mera	sol 176 (training) - onboard	20	9	8	1	5.0%
merb	sol 155 (training) - onboard	20	15	15	0	0.0%
mera	all data - hips	215	109	108	1	0.5%
merb	all data - hips	146	61	60	1	0.7%
mera	all data - onboard	215	111	109	2	0.9%
merb	all data - onboard	146	58	57	1	0.7%

used.

The fhazcams were also carefully calibrated before launch [8]. These models are on-board the rovers, and are used for hazard avoidance and autonomous driving. They were also used for IDD planning and command generation at the beginning of the mission. These models will be referred to as the ATLO models, since the calibration was done in the Assembly, Test, and Launch Operations (ATLO) phase of the project. These models are reported to have mean stereo ranging performance of 3.9mm of a very limited data set during ATLO [6].

Before launch, the IDD was also carefully calibrated, resulting in ATLO IDD models that tested to a mean positioning accuracy of 2.081mm for Spirit and 1.331mm for Opportunity, and a mean repeatability of 0.34mm for Spirit¹[6].

Combining the open-loop² positioning accuracy of the IDD and the stereo ranging performance of the fhazcams, it is expected that before launch the IDD could be positioned to within 4.42 mm for Spirit and 4.12 mm for Opportunity [6], [9].

In Summer 2004, after some time doing operations on the surface, some degradation in the positioning accuracy of the IDD was noticed. By capturing the 3D position of the MBCP from the IDD model and manually capturing the 2D locations of the MBCP images from fhazcam imagery, new camera models were created. This procedure is referred to as the Hybrid Image-Plane/Stereo (HIPS) technique [10]. These models have been in use for IDD planning and command generation since this time.

Once the KVRes tool was verified as described above, much more 5D data (kinematic 3D position and image 2D location) were available. The data from the beginning of the mission

to sol 600 were used with the HIPS technique to generate a new set of HIPS camera models, referred to henceforth as the HIPS2 models.

One use of this analysis tool is to compare these three sets of camera models. The tool can analyze the continuing utility of the HIPS models and expected benefit of the use of the HIPS2 models.

4. RESULTS

As mentioned above, one use of the KVRes tool is to compare the impact of different camera models for the front fhazcams on the Kinematic-Vision Residuals. In this section, we show this comparison for the original ATLO models, for the HIPS camera models, and for the proposed HIPS2 camera models. In these plots, only points that were automatically accepted by the fiducial detector are shown, resulting in different number of data points for the same time span. False positives have been manually removed from these results.

“Early” Data

From egress to approximately sol 160, only low-resolution fhazcam images of the MBCP are available. The three sets of camera models were evaluated based on these images, to investigate the possibility of a post-landing physical camera shift.

Thus, the KVRes tool was used to analyze the Kinematic-Vision Residuals for this time period, with each set of models. Summary plots for this time period are shown in Figures 5 and 6, for MERA (Spirit) and MERB (Opportunity), respectively. The RSS (Root Sum of Squares) length of the 3D vector between the predicted location of the MBCP according to the kinematic model and the detected location of the MBCP as measured by vision is shown on the Y axis, and the time is shown on the X axis. This vector is termed the Kinematic-Visual Residual, or simply the Residual. The 10mm specification for positioning accuracy of the MBCP [1] is shown as a horizontal red line.

¹No repeatability tests were reported for Opportunities IDD.

²with respect to vision - there are closed-loop joint controllers on the IDD.

In Figures 5 and 6, the first row (a) shows the Residual length when the ATLO camera models are used to triangulate the detected position of the MBCP. The second row (b) shows the same length, when the HIPS camera models, and the last row (c) shows the results with the proposed HIPS2 camera models.

For these data, the HIPS models performed the best, with the HIPS2 models performing almost as well. The ATLO models did not result in good agreement between the kinematic predictions of the position of the MBCP and the observed location of the MBCP. However, the precision of these predictions is not as high as the experiment below, as only low-resolution fhazcam images of the MBCP were recorded.

HIPS Training Sol

As described in detail in [10], in Summer 2004 (sol 176 for mera and sol 160 for merb) a set of fhazcam images were taken and manually processed to create the HIPS camera models. This set of images was processed with KVRes, with summaries given in Figure 7 and 8. Some of the training data were rejected by the false-positive rejection tools described above. As before (a) shows the performance with ATLO camera models, (b) with HIPS models, and (c) with HIPS2 models. The HIPS models perform the best on these data. This is to be expected since these are the data with which these models were trained. The HIPS2 models perform almost as well, and the ATLO models perform the least well.


Recent Data

As mentioned above, all but the most recent KVRes results were used to create a new set of HIPS models (the HIPS2) models. In this section, the most recent data (i.e. those not used to create the models) are analyzed for KVRes using each set of models, with summaries given in Figure 9 and 10. These results show that the increased performance of the HIPS2 models in the region of the training data does generalize over the life of the mission, resulting in more accurate predictions of where the MBCP will appear in fhazcam images. More importantly, this characteristic implies better positioning accuracy of the IDD on science targets.


Full Data Set

Finally, the KVRes tools is run on all available high-resolution data, with summaries given in Figure 11 and 12. Before approximately sol 160, only low-resolution data were taken from the fhazcams. To avoid problems comparing the low-resolution and high-resolution camera models, only the latter are used.


These data illustrate the performance of the various camera models over the complete data set. The mean residual is reduced from 6.11mm and 3.82mm using HIPS to 3.13mm and 2.77mm, for MERA (Spirit) and MERB (Opportunity) respectively, using HIPS2. The median residual is reduced from 6.28mm and 3.01mm to 2.83mm and 2.01mm. Finally,



(a) ATLO Models

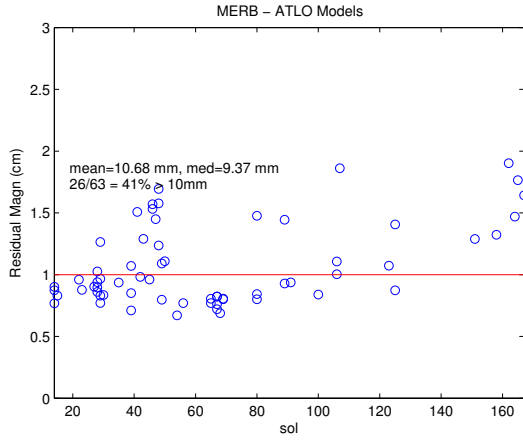


(b) HIPS Models




(c) HIPS2 Models


Figure 5. KVRes for all models, on early data (sol 0 to 160), for MERA (Spirit)



(a) ATLO Models

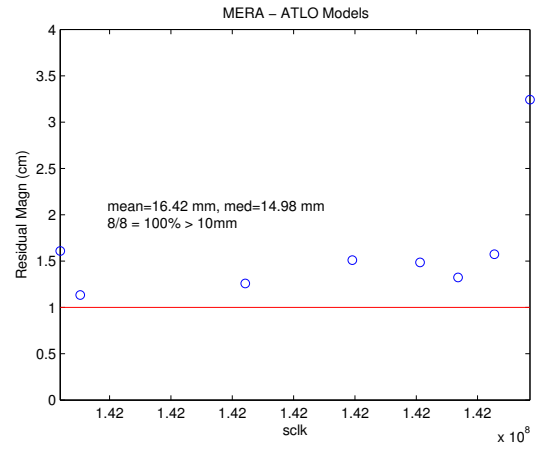


(b) HIPS Models

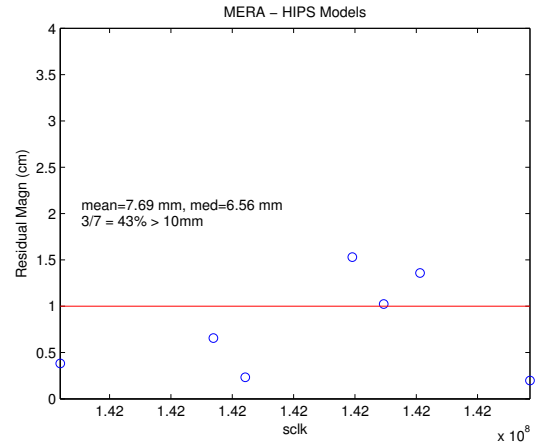


(c) HIPS2 Models

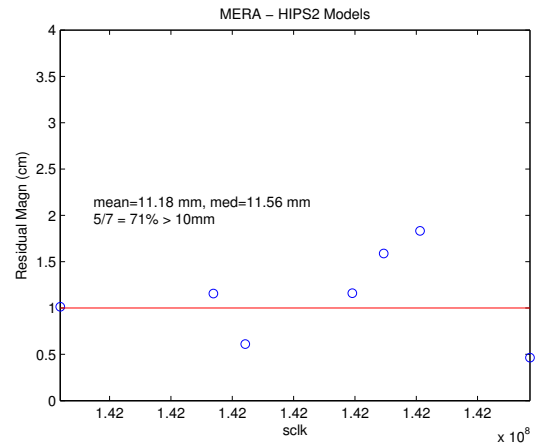
Figure 6. KVRes for all models, on early data (sol 0 to 160), for MERB (Opportunity)



(a) ATLO Models




(b) HIPS Models




(c) HIPS2 Models

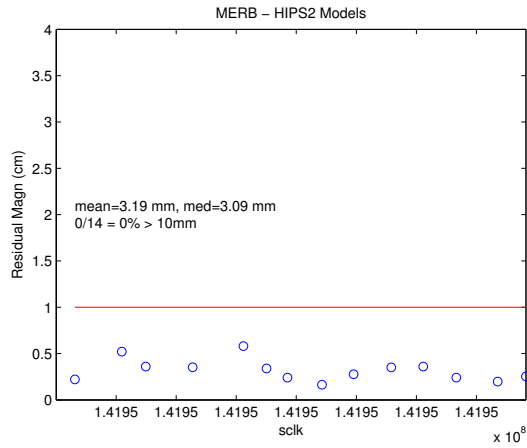
Figure 7. KVRes for all models, on HIPS training data (sol 155), for MERA (Spirit). The x axis is time: all images are from the same sol.



(a) ATLO Models




(b) HIPS Models

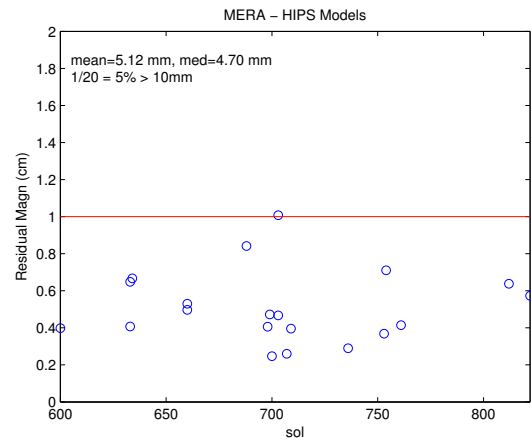


(c) HIPS2 Models


Figure 8. KVRes for all models, on HIPS training data (sol 176), for MERB (Opportunity). The x axis is time: all images are from the same sol.



(a) ATLO Models

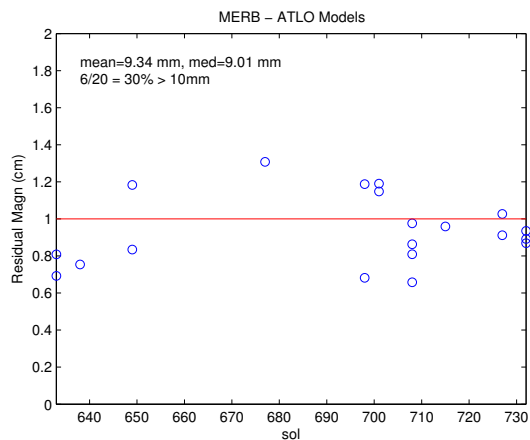


(b) HIPS Models

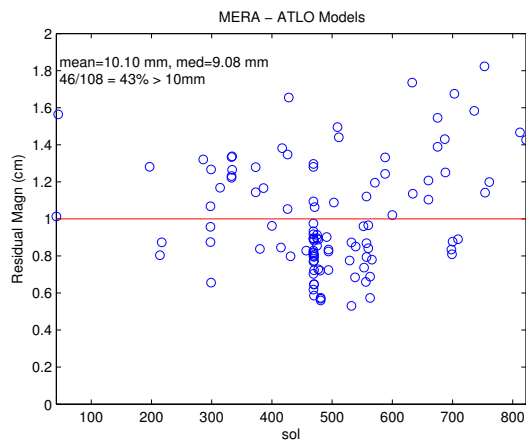


(c) HIPS2 Models

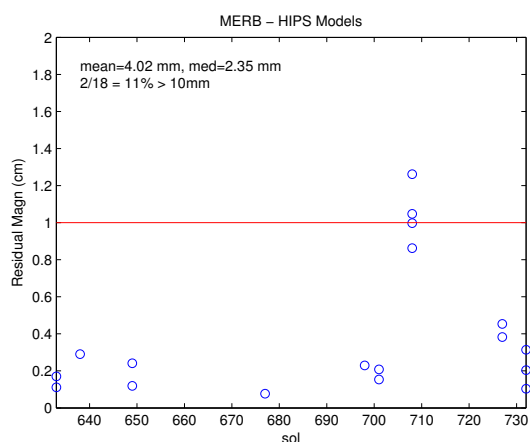
Figure 9. KVRes for all models, on recent data (since sol 600), for MERA (Spirit)



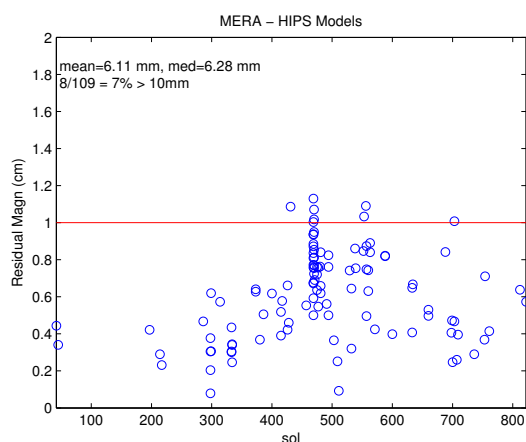
(a) ATLO Models



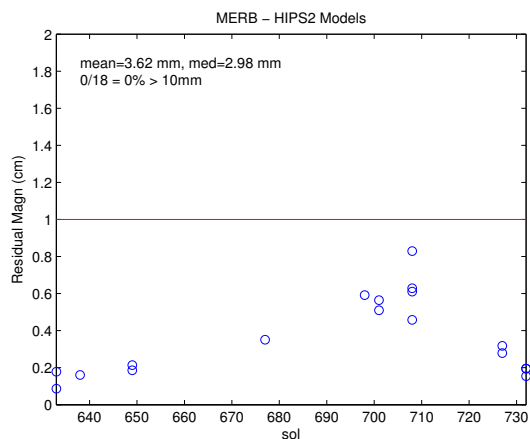
(a) ATLO Models




(b) HIPS Models



(b) HIPS Models



(c) HIPS2 Models



(c) HIPS2 Models

Figure 10. KVRes for all models, on recent data (since sol 600), for MERB (Opportunity)

Figure 11. KVRes for all models, on all high resolution data (since sol 160), for MERA (Spirit)

the number of predictions above the 10mm threshold is reduced from 7% and 4% to 1% and 0%.

5. CONCLUSIONS

The KVRes tool was developed to track the length of the residual between the kinematic prediction of the position of the Mössbauer Contact Plate and the location as detected by vision. This residual gives some indication of the health of the complete Vision-IDD system, which impacts the positioning accuracy of the IDD. The positioning accuracy of the IDD is an important characteristic for the collection and analysis of science data.

This analysis *does not reveal* the distribution of the errors between the vision and IDD systems, as it only compares the two predictions for the Mössbauer Contact Plate position. Some indications, such as continued nominal stow/unstow characteristics of the IDD, point to the majority of this error being camera-related, but this is an open discussion at press time.

The pre-flight camera models, as well as two sets of camera models developed since the vehicles have landed on Mars, have been compared with the length of the KV Residual used as a metric. This illustrates that the HIPS models that have been in use since Summer 2004 yield significant improvements in prediction accuracy, and that a proposed set of HIPS2 models would yield some additional accuracy in predictions.

6. FUTURE WORK

Some method for automatically flagging suspected false positives for manual review should be implemented. Even a simple measure such as enforcing the review of all KVRes measurements above the 10mm design specification would reduce the workload to a very manageable size. For example, 8 of the 109 image pairs on MERA (Spirit) would need to be reviewed, and 2 of 57 on MERB (Opportunity).

The current models locate the position of a single point on the IDD tool turret, and track this point. It would be desirable to additionally detect and track the *orientation* of the IDD, and compare this against the design specification of 5° in orientation accuracy for novel science targets.

Generating camera models based on different time periods and comparing KV Residuals using these camera models may reveal more information about if and how the closed-loop Vision-IDD system is evolving over time.

ACKNOWLEDGMENT

This work was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration. The work involves important contributions from many colleagues at

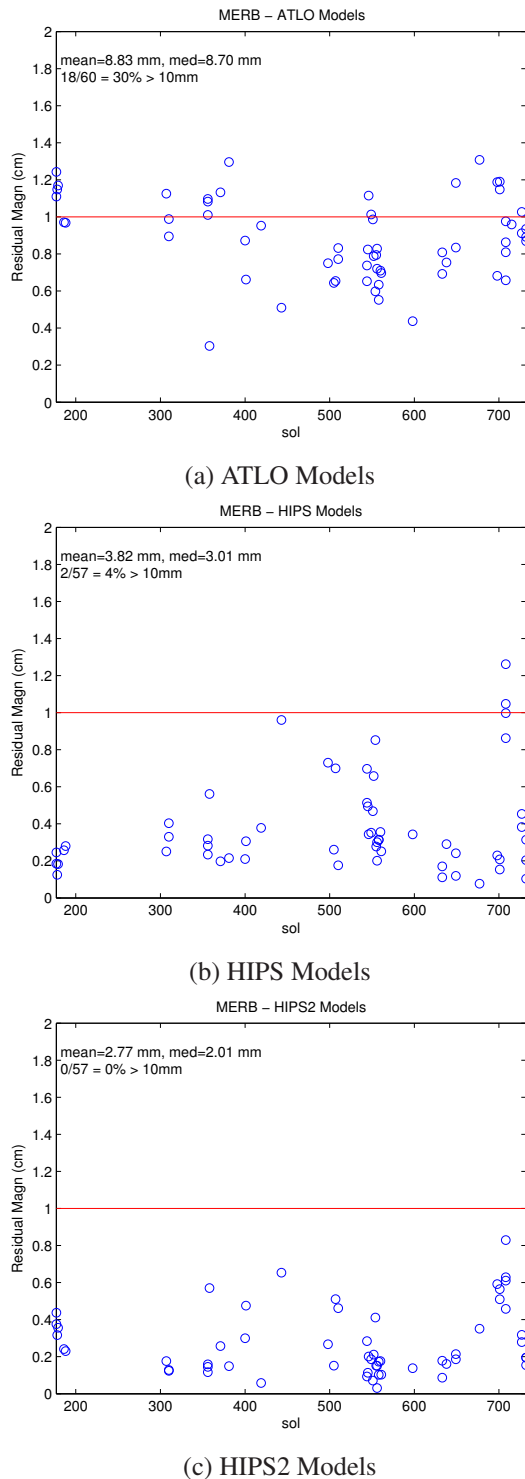


Figure 12. KVRes for all models, on all high resolution data (since sol 160), for MERB (Opportunity)

both JPL and collaborating institutions. We gratefully acknowledge these interactions and note many of the specific developments in references that follow. Some of the background of planetary manipulation appears in [10].

This work was carried out while Dr. Nickels was on an academic leave from Trinity University. Trinity's contribution to his stay at the Jet Propulsion Laboratory during the 2005-2006 academic year is gratefully acknowledged.

REFERENCES

- [1] S. Squyres *et al.*, "Athena Mars rover science investigation," *Journal of Geophysical Research*, vol. 108, no. E12, p. 8062, Nov. 2003.
- [2] J. Erickson *et al.*, "Mars Exploration Rover surface operations," in *Proceedings of the 53rd International Astronautical Congress, The World Space Congress*, Houston, TX, Oct. 2002, paper No. IAC-02-Q.3.1.03.
- [3] J. Shi and C. Tomisito, "Good features to track," in *Proceedings IEEE Computer Society Conference on Computer Vision Pattern Recognition*, June 1994, pp. 593–600.
- [4] J. P. Lewis, "Fast normalized cross-correlation," in *Vision Interface*, 1995.
- [5] D. B. Gennery, "Generalized camera calibration including fish-eye lenses," *International Journal of Computer Vision*, 2006, to appear.
- [6] E. T. Baumgartner *et al.*, "The Mars Exploration Rover instrument positioning system," in *Proc. IEEE Aerospace Conference*, Big Sky, MT, Mar. 2005.
- [7] E. T. Baumgartner, R. G. Bonitz, J. P. Melko, L. R. Shiraishi, P. C. Leger, and A. Trebi-Ollennu, "Mobile manipulation for the Mars Exploration Rover," *IEEE Robot. Automat. Mag.*, vol. 13, no. 2, pp. 27–36, June 2006.
- [8] J. N. Maki *et al.*, "Mars Exploration Rover engineering cameras," *Journal of Geophysical Research*, vol. 108, no. E12, p. 8071, Nov. 2003.
- [9] E. T. Baumgartner, "MER IPS accuracy," Personal Correspondance, Nov. 2006.
- [10] M. Robinson, E. Baumgartner, T. Litwin, and K. Nickels, "Hybrid image plane/stereo (HIPS) manipulation for robotic space applications," , Under Review.



Kevin Nickels is an associate professor in the Department of Engineering Science at Trinity University. He received the B.S. degree in Computer and Electrical Engineering from Purdue University (1993), and received the M.S. degree (1996) and the Ph. D. (1998) in Electrical Engineering from The University of Illinois at Urbana- Champaign. He is currently working in the areas of computer vision, pattern recognition, and robotics. Dr. Nickels spent the 2005-2006 academic year as a Senior Engineer in Avionics Systems and Space Technology at NASA's Jet Propulsion Laboratory, concentrating on Vision Guided Manipulation.



Max Bajracharya is a member of the Mobility and Manipulation Group at the Jet Propulsion Laboratory, Pasadena, CA. His research includes vision-based tracking, manipulation, navigation, and pose estimation for Mars rovers. He is currently the task lead for several tasks focusing on autonomous target approach and instrument placement. Max received his Bachelors and Masters degrees in computer science and electrical engineering from MIT in 2001.



Ashitey Trebi-Ollennu Ashitey Trebi-Ollennu is a Senior Member of Engineering Staff at the NASA-Jet Propulsion Laboratory, California Institute of Technology, in Pasadena, CA and a Fellow of the IEE (U.K.). His current research at JPL focuses on Planetary Rovers, Manipulation, Multiple Mobile Robots (Planetary Outpost), Reconfigurable Robots and Man-machine Interaction. Dr. Trebi-Ollennu's research has resulted in more than 65 publications. He currently works on the Mars Exploration Rover Operations Team as a Rover Planner (Rover Driver) responsible for Surface Mobility/Navigation Planning, IDD Planning, and Command Generation.



Robert Liebersbach is a member of the Mobility and Manipulation Group at the Jet Propulsion Laboratory, Pasadena, CA. He is currently an undergraduate Mechanical Engineer at the University of California, Irvine. His area of research includes long term mission mapping, autonomous health assessment, and calibration. Robert currently works on the Mobility and IDD subsystem of the Mars Exploration Rover mission.