

July 2008

Autonomous geometric precision error estimation in low-level computer vision tasks

Andrés Corrada-Emmanuel

University of Massachusetts - Amherst, corrada@cs.umass.edu

Howard Schultz

University of Massachusetts at Amherst, hschultz@cs.umass.edu

Follow this and additional works at: http://scholarworks.umass.edu/cs_faculty_pubs



Part of the [Computer Sciences Commons](#)

Recommended Citation

Corrada-Emmanuel, Andrés and Schultz, Howard, "Autonomous geometric precision error estimation in low-level computer vision tasks" (2008). *International Conference on Machine Learning 2008*. 2.

http://scholarworks.umass.edu/cs_faculty_pubs/2

This Article is brought to you for free and open access by the Computer Science at ScholarWorks@UMass Amherst. It has been accepted for inclusion in Computer Science Department Faculty Publication Series by an authorized administrator of ScholarWorks@UMass Amherst. For more information, please contact scholarworks@library.umass.edu.

Autonomous Geometric Precision Error Estimation in Low-Level Computer Vision Tasks

Andrés Corrada-Emmanuel
Howard Schultz

CORRADA@CS.UMASS.EDU
HSCULTZ@CS.UMASS.EDU

Computer Science Department, University of Massachusetts at Amherst, 140 Governors Drive, Amherst, MA 01003-9264

Abstract

Errors in map-making tasks using computer vision are sparse. We demonstrate this by considering the construction of digital elevation models that employ stereo matching algorithms to triangulate real-world points. This sparsity, coupled with a geometric theory of errors recently developed by the authors, allows for autonomous agents to calculate their own precision independently of ground truth. We connect these developments with recent advances in the mathematics of sparse signal reconstruction or compressed sensing. The theory presented here extends the autonomy of 3-D model reconstructions discovered in the 1990s to their errors.

1. Introduction

Autonomy of robots or intelligent sensors depends on developing algorithms that can assess their own performance independent of ground truth. Consider an Aerial Mapping Appliance (AMA) that must construct a map or 3-D model of the world based on photographs. Researchers in computer vision discovered in the 1990s that a faithful 3-D model of an imaged scene was possible without any knowledge of the positions, or orientations of the camera that took the photographs (Beardsley et al., 1996). This reconstruction is even possible without knowing the internal parameters of the camera (Pollefeys et al., 1999). The geometry of multiple images contains all the necessary information to do this reconstruction. This independence of 3-D model reconstruction from ground truth

(in this case, camera positions, etc.) raises the possibility that the errors in the reconstruction can also be recovered autonomously by an intelligent agent such as the AMA.

Autonomous error estimation for 3-D model reconstruction was recently demonstrated to be possible by the authors (Corrada-Emmanuel et al., 2007; Corrada-Emmanuel & Schultz, 2008). The theory depends on making a distinction between accuracy and precision. Knowledge of accuracy is not possible without ground truth. Precision can be estimated autonomously. This paper will demonstrate that autonomous precision estimation is also related to the mathematics of sparse signal reconstruction or compressed sensing (Donoho, 2006a). The precision errors of measurements, not just the measurements, are sparse themselves. This sparsity is the key to their reconstruction.

2. The Distinction between Geometric Accuracy and Precision

The concepts of accuracy and precision are well known to all scientists. The Machine Learning community knows these concepts as bias and variance (Bishop, 2007). Bias refers to how far an estimate is from the true value. Variance captures how noisy that estimate is given the measurements used to compute it. Our meaning of accuracy and precision in 3-D models is analogous to bias and variance but not equivalent. Our definitions are geometrical in nature.

Imagine that one had a set of 3-D models of a scene. Furthermore, along with the models one also has the ground truth or exact locations of points in the scene. The total error of the models can be defined as

$$\sum_{\text{models}} \sum_{\text{points}} [\vec{x}_{\text{model}}(\text{point}) - \vec{x}_{\text{true}}(\text{point})]^2 \quad (1)$$

We can decrease the total error in the models by applying a global transformation to all the models. For

Appearing in *Proceedings of the 25th International Conference on Machine Learning*, Helsinki, Finland, 2008. Copyright 2008 by the author(s)/owner(s).

example, the models may be systematically off by 1 meter in some direction. The *geometric precision* of the models is defined as the minimum possible total error after application of a global transformation to all the models:

$$\min_T \sum_{\text{models}} \sum_{\text{points}} [T(\vec{x}_{\text{model}}(\text{point})) - \vec{x}_{\text{true}}(\text{point})]^2. \quad (2)$$

The *geometric accuracy* is defined as difference between the total error and the geometric precision.

We describe some simple examples to clarify these definitions. Imagine a 3-D model reconstruction that is merely a translated example of the real world, i.e. all locations are off by 1 meter to the North. The reconstruction has an accuracy error of 1 meter and zero precision error. A model with zero accuracy error but some precision error can be created by taking a perfect reconstruction and individually randomizing the elevation of the reconstruction with zero mean.

The concept of geometric accuracy can thus be captured by a successive set of transformations that the reader may want to view as encompassing the usual hierarchy of projective, affine and euclidean transformations (Hartley & Zisserman, 2000; Faugeras et al., 2001) or some subset of them. The rest of this paper will use the words accuracy and precision as shorthand for these geometric definitions.

2.1. Autonomous Elevation Difference Equations and Geometric Precision

Our central claim is that precision can be estimated autonomously even as the accuracy of the models is completely unknown. Autonomous geometric precision error estimation is possible by creating quantities that are invariant under global accuracy transformations. In this paper we will consider one such quantity that is useful for characterizing the precision errors in DEMs and has been discussed in our previous papers (Corrada-Emmanuel et al., 2007):

$$\Delta_{P,Q}(x, y) = \frac{1}{P} \sum_{i=1}^P Z_i - \frac{1}{Q} \sum_{j=1}^Q Z_j \quad (3)$$

$$= \frac{1}{P} \sum_{i=1}^P \delta_i - \frac{1}{Q} \sum_{j=1}^Q \delta_j, \quad (4)$$

where the integers P and Q are between 1 and the number of models being compared, $1 \leq P \leq M$ and $1 \leq Q \leq M$. A DEM i is a collection of elevation postings at different (x, y) locations, $\{Z_i(x, y)\}$. The precision error in each posting is denoted by $\delta_i(x, y)$,

so in general one can write

$$Z_i(x, y) = Z_{\text{true}}(x, y) + \delta_i(x, y). \quad (5)$$

By picking the integers P and Q less than or equal to the number of DEM models, we guarantee that the true value of the elevation cancels out at each posting since

$$\frac{1}{P} * (P * Z_{\text{true}}) - \frac{1}{Q} * (Q * Z_{\text{true}}) = 0 \quad (6)$$

so that equation 4 follows from equation 3.

By considering all possible values P and Q one can find a set of linearly independent equations for the elevation precision errors. We call this independent set the *autonomous difference equations*. Note that these equations are not being used to construct a better estimate of the true elevation by performing some simple averaging over them. Their sole purpose is to probe the errors in the reconstructed elevations. More general expressions that take into account x and y position errors can be constructed but we defer discussion of these to future papers.

Equation 3 can be calculated from the observable elevations. The task of the autonomous agent is to estimate the precision errors $\{\delta_i\}$ in equation 4 and how they are correlated with each other. Once the agent knows these correlations, the precision error of a fused estimated can be decreased while possibly increasing its accuracy error. This may be a suitable action to take since in many computer vision tasks accuracy is cheaper to fix than precision, a point we clarify in our concluding remarks.

3. The Covariance Matrix for Precision Errors

An AMA or robot on a mapping mission will not know beforehand what errors it will make during its activities. Sensors could systematically malfunction. Lighting conditions may be unfavorable at certain viewing angles. These and other factors will inevitably mean that repeated measurements of the same scene will be partly correlated, or their precision may vary widely. How should the 3-D models obtained from different vantage points be fused? How fast is the precision error in the reconstructions decreasing as a function of the collected images? Has the AMA attained a desired precision level and therefore completed its mission? Autonomous mission planning by robots requires answers to these questions.

We argue that a principled approach to answering these questions must rely on an autonomous estima-

tion of the covariance matrix of the 3-D models. Having multiple measurements whose errors are strongly correlated is not much better than a single measurement, for example. Knowing the covariance matrix would allow the agent to discard bad data, understand its rate of error decrease as a function of data collection, and provide a fused estimate that monotonically improves with time.

The covariance matrix for DEM errors is composed of entries of the form $\langle \delta_i \delta_j \rangle - \langle \delta_i \rangle \langle \delta_j \rangle$. For ease of discussion, we will assume that the precision error has been de-measured so $\langle \delta_i \rangle = 0$ for all i , so the covariance matrix is equivalent to $\langle \delta_i \delta_j \rangle$ in this paper.

It is impossible, generally, to calculate this covariance matrix given a set of measurements. We explain this fully by constructing a linear algebra system for the covariance matrix based on the autonomous difference equations (eqs. 3 and 4) to demonstrate that it defines an under-determined system – one where we have less equations than unknowns.

3.1. An Under-Determined Linear Algebra System for the Covariance Matrix Entries

Squaring the autonomous difference equations and averaging over all the posting locations (x, y) gives a set of linear equations for all the entries in the covariance matrix. We denote this system by

$$S = \Phi \Delta. \quad (7)$$

The vector S is the “signal” of the DEM precision errors. Its components are calculated using equation 3. The matrix Φ consists of the rational fractions that come from expanding the square of equation 4. The vector Δ are the entries $\langle \delta_i \delta_j \rangle$ of the covariance matrix that we want to estimate.

Equation 7 defines an under-determined linear system because the number of independent entries in the covariance matrix is $M(M + 1)/2$ given M models (the matrix is symmetric). The number of independent equations that can be constructed from the autonomous difference equations is equal to $M(M + 1)/2 - M$. Therefore, the system is always under-determined by M equations.

3.2. The Correlated-Pair Error Model

This limitation was circumvented in our earlier papers (Corrada-Emmanuel et al., 2007) by *assuming* that the

covariance matrix had the simple form

$$\begin{pmatrix} * & * & 0 & 0 & \dots \\ * & * & 0 & 0 & \dots \\ 0 & 0 & * & * & \dots \\ 0 & 0 & * & * & \dots \\ \vdots & \vdots & \vdots & \vdots & \ddots \end{pmatrix} \quad (8)$$

The block-diagonal shape came from our production of two DEMs from every photographic pair, a practice that differs from the usual photogrammetric convention of producing a single DEM from a photographic pair. We assumed that only the two DEMs from the same photographic pair were correlated with each other. These correlated-pair DEMs gave rise to the block-diagonal form of the covariance matrix. In effect, we assumed that the covariance matrix was sparse since this correlated-pair model only requires $n + n/2$ non-zero terms to be estimated for the covariance matrix.

3.2.1. ASYMMETRY IN COMPUTER VISION STEREO MATCHING

The reason one can produce two different DEMs from two photographs is that stereo matching algorithms may not be perfectly symmetric in their output (Brown et al., 2003). This means that a DEM produced by matching image A to image B, which we denote by $A \rightarrow B$, will not lead to the same DEM as doing $B \rightarrow A$. Of course, the resulting DEM pair ($A \rightarrow B, B \rightarrow A$) is highly correlated. The correlated-pair error model in equation 8 is meant to capture this unknown, but possibly large, cross-correlation between the errors in the pair. One can readily calculate that this block-diagonal error model allows one to calculate the precision error exactly whenever three or more photographs overlap on the same scene.

4. Sparsity of Geometric Precision Error

As useful as the correlated-pair error model may be in certain circumstances it is a *model* and therefore cannot form the foundation of a robust process for error estimation. It is conceivable that DEMs from unrelated photographs could become correlated in their errors due to environmental factors or even instrument malfunction. A robust estimation of the covariance matrix should not depend on any assumptions of how DEMs are correlated.

Recent developments in the mathematics of sparse signal reconstruction or compressed sensing (Donoho, 2006a) offer us a mathematical procedure to deal with

this situation. During a mapping mission photographs taken from different viewing positions and orientations will lead to mapping errors that are uncorrelated but occasionally may have strong correlations between them. We just do not know a priori which DEMs will be correlated with each other, only that these cross-correlations will be sparse.

As we noted in section 3.1, the linear system is at the margin of being completely determined being shy by just M equations. If the covariance matrix was sparse enough in the sense that on the order of M independent entries were zero, the estimation would be robust. We can express this condition by dividing the number of equations we have $(M(M+1)/2 - M)$ by the number we need $(M(M+1)/2)$

$$1 - \frac{2}{M+1}. \quad (9)$$

As the number of models increases, this fraction becomes increasingly near to one – the condition for being well-determined.

We hypothesize that given enough models ($M \rightarrow \infty$), any experimental situation can be driven into a sparse regime for the precision error covariance matrix. Note that we are not talking about sparsity of the models themselves, but of the correlations between their precision errors.

The under-determined linear system 7 can be solved by using the ℓ_1 -minimization technique advocated in the compressed sensing literature (Donoho, 2006b)

$$\min \|\Delta\|_1 \text{ subject to } S = \Phi\Delta. \quad (10)$$

This problem can be solved as a convex optimization problem (Donoho, 2006b) by recasting it as the equivalent linear program:

$$\min \sum_i u_i \text{ subject to} \quad (11)$$

$$u_i + \Delta_i \geq 0 \quad (12)$$

$$u_i - \Delta_i \geq 0 \quad (13)$$

$$\Phi\Delta = S \quad (14)$$

In the experimental section of the paper we will show that this approach reconstructs a covariance matrix for the precision errors that is very close to the correlated-pair model (eq. 8). Some off-diagonal terms hypothesized to be zero are about 5 times smaller than the in-pair cross correlation. We emphasize that this sparsity of errors hypothesis is an experimental assertion. No mathematical proof can be given that this sparsity condition can be met. The applicability of the

assertion is based on the experimental realization of the AMA and the features of the terrain. A device built with stable imaging sensors of high quality that is mapping a reasonably static terrain would be a good candidate for a suitable condition that meets our sparsity assumption.

5. More Data Means Higher Resolution Error Maps

The name *compressed sensing* comes from the realization that sparsity implies a low-dimensional or compressible signal. If pictures of a natural scene taken with a $n \times n$ CCD can always be compressed, why take n^2 measurements? The imaging of the scene can be compressed by using less pixels and then reconstructed with an under-determined linear system. This has been dramatically demonstrated by the Rice University one-pixel camera (Wakin et al., 2006). About 1,000 measurements with a single pixel reproduced images captured by a 4,000 pixel CCD. Compressed sensing implies that we are wasting effort by taking too many measurements.

The error theory presented here gives a different perspective on this issue. Yes, reconstructing a 3-D model of the world can be done with less measurements. However, errors are an important aspect of all measurements. How confident can we be of any particular reconstruction? The only way to understand this is to produce not just maps of the territory that is being mapped, i.e. DEMs, but to also produce error maps of the same territory. The procedure for precision error estimation depends on averaging over all postings that a collection of DEMs have in common. Sparsity is only present after this averaging. The error map therefore has a much lower resolution than the DEM itself. Multiple measurements are needed to increase the resolution of this error map. In this view, no measurement is ever wasted – it leads to higher resolution in the error map of the measurements.

This suggests that the resolution of the error map should be studied by decreasing the map area that is used to create the average covariance matrix of the precision errors. As the averaging area is diminished, various cross-correlations between different DEMs will start to turn on. At some point, the number of these off-diagonal terms will be large enough to violate the condition of sparsity and the resolution limit of the error map would be reached. This resolution limit may vary across the mapped area and would naturally depend on the particular dataset. This phenomenon will be demonstrated in the experimental section.

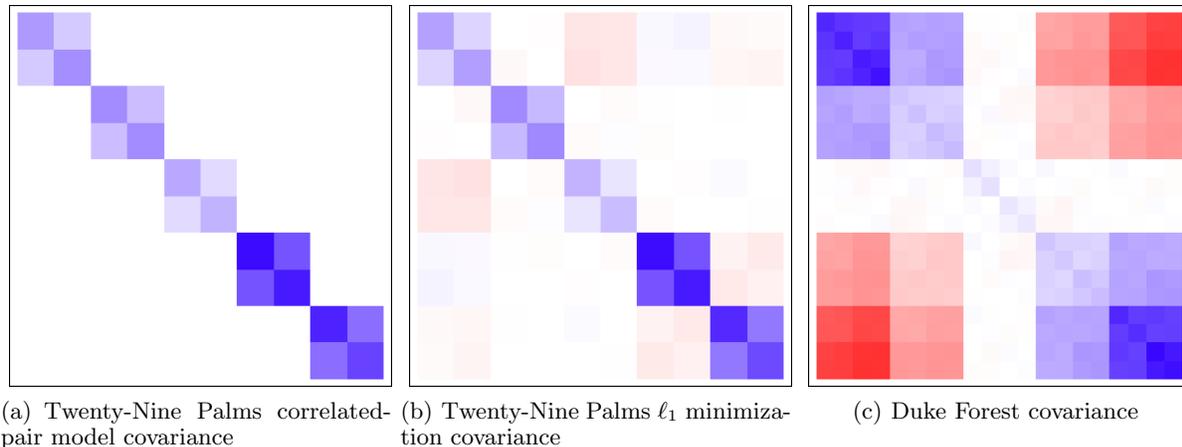


Figure 1. Comparison of the covariance matrix obtained with ℓ_1 -minimization versus that obtained by using the correlated-pair error model. All values are in units of m^2 . Blue represents positive values, red negative values. The hue scale for figures (a) and (b) is normalized so that a value of $0.12 m^2$ results in a completely saturated color pixel. Figure (c) is normalized with $2.5 m^2$.

6. Experimental Results

We demonstrate the formalism for sparse precision error estimation by using a set of four aerial images taken of a desert terrain in the Twenty-Nine Palms region in California, USA. The images have been arbitrarily labeled as $\{A, B, C, D\}$. Four photographs allow us to produce $4 \times 3 = 12$ DEMs from all possible *matching chains* of the form $i \rightarrow j$. A blunder removal process, however, automatically identified that two of the DEMs ($B \rightarrow D$, and $D \rightarrow B$) differed in their elevation estimates by more than one meter for all postings. This pair was excluded from our calculations so the results presented here involve the remaining 10 DEMs.

6.1. Random Reconstructions via ℓ_1 -Minimization

For 10 DEMs there are on the order of fifty thousand ways to write equation 4. The number of independent equations in this set is 45. An independent set selected from all possible permutations of the difference equations leads to a different reconstruction matrix Φ . To carry out the ℓ_1 -minimization estimate of the 10×10 covariance matrix for the DEMs, we randomly selected ten different linearly independent sets and their corresponding Φ matrices. This was done to study the numerical stability of the reconstruction procedure. The statistical average was done over an overlap region of the DEMS that spanned the postings 500 to 1500 where 2000 by 2000 was the original size of the individual DEMs. This was done to exclude edge effects and increase the density of postings on which all DEMs gave an elevation estimate. Only postings

for which we had a full 10 measurements were used. The number of postings was equal to 940,010 out of a possible million. Each posting represents an area of $(0.38 m)^2$.

The reconstruction of the covariance matrix is shown in figure 1(b). The covariance matrix is presented as a 10×10 pixel image. For comparison, the covariance matrix reconstructed with the correlated-pair error model is shown in figure 1(a). No numerically significant variation in the reconstructed covariance matrix was observed with 10 randomly selected Φ matrices so a single figure is sufficient to summarize the results. Note that about 12 entries in the covariance matrix are practically zero – two more than the 10 entries required to define a well-determined linear system.

The ℓ_1 -minimization procedure also ascribes most of the cross-correlations to the DEMs that come from asymmetrically matching the same pair of photographs. In addition, the sparse reconstruction has discovered that some of the DEMs are negatively correlated.

Is this error reconstruction correct? At this time we can point to its *self-consistent* character as strong evidence for its correctness. Neither the autonomous elevation difference equations or the ℓ_1 -minimization procedure assume that certain DEMs are strongly correlated. Yet the empirically reconstructed matrix clearly shows that the 10 DEMs have a strong 5-pair structure exemplified by the block-diagonal structure. The reconstruction has ‘discovered’ that we used DEMs from asymmetric matching pairs.

Another self-consistent feature of the reconstruction is

that the more precise a DEM is, the smaller its cross-correlation with its asymmetric pair becomes. This is a behavior that we would expect from a system that is producing increasingly precise estimates.

6.1.1. DUKE FOREST DEMS

The Twenty-Nine Palms data, just discussed, is extremely high quality. The images were taken with a high-quality photogrammetric instrument. To confirm that precision errors can be recovered in more noisy data, we studied a series of aerial photographs of a forest canopy in the Duke Forest, NC taken with an off-the-shelf digital camera.

We randomly selected a track of images and picked four consecutive images. The images were multi-band and we chose the near-ir and green bands. The combination of bands and asymmetric pair matches resulted in twenty DEMs. The recovered precision error covariance matrix is shown in Figure 1(c). The recovery is similar to that for the Twenty-Nine Palms data, in that the highly-correlated pairs were discovered once again. As befits the noisier data set, the precision error estimated is $\approx 2.0 m^2$ versus the $0.1 m^2$ value for the Twenty-Nine Palms images.

6.2. Horizontal Resolution of the Precision Error Covariance Matrix

The self-consistent character of the reconstruction can be exploited further. The linear program that recovers the error has as side constraints that the diagonal terms of the covariance matrix have to be positive, a required property for the $\langle \delta_i^2 \rangle$ terms. To keep the reconstruction as a linear program, we did not require that the cross-correlations terms satisfy the inequality

$$\left| \frac{\langle \delta_i \delta_j \rangle}{\sqrt{\langle \delta_i^2 \rangle \langle \delta_j^2 \rangle}} \right| \leq 1. \quad (15)$$

The output of the linear program just turns out to satisfy these constraints for this particular dataset. Indeed, we now use the breakdown in these cross-correlation constraints to probe how much resolution can be obtained in the error map of the DEMs.

To study the resolution limit of ℓ_1 -minimization procedure, we shrank the size of the area in the Twenty-Nine maps over which the difference equations were averaged. We have no independent way of verifying the validity of the reconstruction except the self-consistency check that the reconstructed vector does indeed represent a covariance matrix – its dimensionless cross-correlations should have an absolute value less than or equal to one.

Surprisingly the resolution of the covariance matrix for this data is on the order of 5x5 postings. We show an example of the covariance matrix for a patch encompassing the postings 500 to 505 in both directions in figure 2(a). The covariance matrix for the patch encompassing the postings 500 to 510 is shown in figure 2(b). The breakdown in reconstruction for the 5x5 patch is most evident in the cross-correlations related to the DEM in position 9. For this particular patch the variance of DEM 9 is calculated as $1.1 \cdot 10^{-5} m^2$ and the variance of DEM 1 is calculated as $1.0 \cdot 10^{-1} m^2$. The dimensionless cross-correlation between them has a value of 30.0 – the ℓ_1 -minimization has not produced a proper covariance matrix for this small patch.

No breakdown is found in the 10x10 patch. Similar results were obtained over a handful of other patches over the mapped scene. Interestingly, an earlier paper by the authors had found that for this dataset the horizontal decorrelation length was in the order of 5 postings, a result that was obtained by a “cheating” experiment that used ray-tracing to establish a pseudo ground truth against which the error at the individual posting level could be calculated.

7. Conclusions and Future Work

We conclude by discussing the utility of estimating precision error even while neglecting or increasing the accuracy error. Geometric accuracy is defined by a global set of transformations. The parameters needed to define it are finite and readily extracted by knowing at most the location for three points in the world. In that sense, accuracy is cheap to obtain. Precision, on the other hand, captures the local variability of the 3-D model reconstructions. The parameters needed to model it, if one wished to do so, are correspondingly large. Therefore, 3-D model precision is expensive for the user to correct since it involves multiple measurements spread over the whole scene. Therefore, the autonomous error estimation algorithm presented here should have application in computer vision tasks where accuracy is not needed. An example of such a task is species identification by shape where the resolution is more important than the absolute size or orientation of the objects. Many more examples can be thought of, where accuracy is not relevant but precision is. One obvious class of problems for which this algorithm would not be helpful is those that require accurate geolocation of an object in the scene. A reasonable guarantee of accuracy can be obtained by using proper external references (GPS, attitude-heading sensors, etc.) but this algorithm is invariant to their accuracy error.

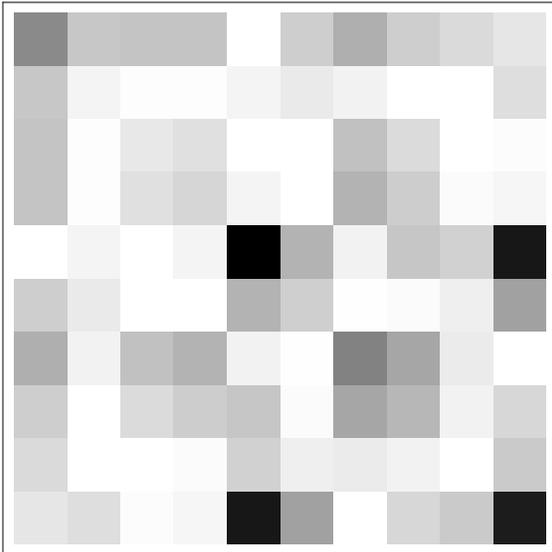
Table 1. Covariance matrix entries along the block diagonal for the ten DEMs in the 29 Palms dataset. Variance is in units of m^2 .

DEM	$\langle \delta_i^2 \rangle$	$\langle \delta_i \delta_j \rangle / \sqrt{\langle \delta_i^2 \rangle \langle \delta_j^2 \rangle}$
AB	0.044	0.45
BA	0.046	
AC	0.056	0.59
CA	0.056	
AD	0.036	0.38
DA	0.031	
BC	0.114	0.73
CB	0.108	
CD	0.100	0.69
DC	0.085	

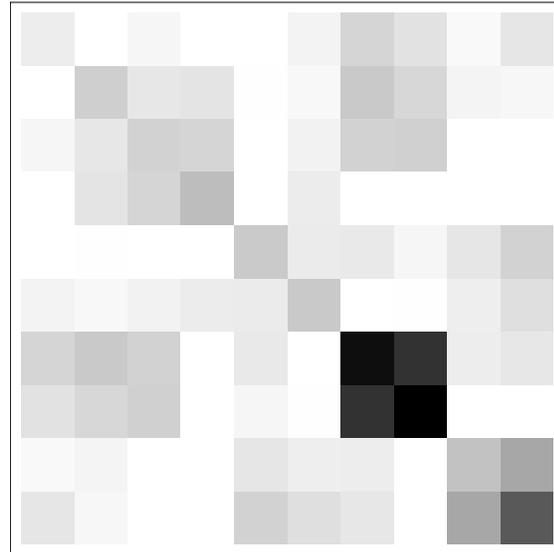
(a) ℓ_1 -minimization

DEM	$\langle \delta_i^2 \rangle$	$\langle \delta_i \delta_j \rangle / \sqrt{\langle \delta_i^2 \rangle \langle \delta_j^2 \rangle}$
AB	0.048	0.50
BA	0.053	
AC	0.054	0.57
CA	0.054	
AD	0.041	0.44
DA	0.036	
BC	0.115	0.73
CB	0.108	
CD	0.104	0.71
DC	0.089	

(b) Correlated-pair error model



(a) Covariance matrix for a patch of size 5x5



(b) Covariance matrix for a patch of size 10x10

The present paper has demonstrated that the covariance matrix of the geometric precision errors can be measured autonomously. The present formalism applies to other areas of Machine Learning, indeed, to any scientific setting where multiple scalar predictions are available for a set of entities. For example, the precision error equations (3) can be used for comparing the relevance judgment of various information retrieval algorithms. Instead of elevations, one would use the binary judgment of relevancy to compare the retrieval models.

Future work will continue to explore the utility of the precision error covariance matrix for data fusion (what is the optimal way to combine the DEMs to minimize the total precision error?), and extend the formalism of the precision error equations (3) to multi-dimensional or other non-scalar models such as 3-D locations or parse trees in computational linguistics.

Acknowledgments

This work was supported by grant (IIS-0430742) from the United States National Science Foundation and contract TT0690688 from the Advanced Technology Laboratory, Lockheed-Martin Corporation.

References

- Beardsley, P., Torr, P., & Zisserman, A. (1996). 3D model acquisition from extended image sequences. *Proceedings of the 4th European Conference on Computer Vision, ECCV'96* (p. 683). Cambridge, UK: Springer-Verlag. Part 2 (of 2).
- Bishop, C. M. (2007). *Pattern recognition and machine learning*. Springer.
- Brown, M. Z., Burschka, D., & Hager, G. D. (2003). Advances in computational stereo. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 25, 993–1008. Compilation and indexing terms, Copyright 2007 Elsevier Inc. All rights reserved.
- Corrada-Emmanuel, A., Pinette, B., Ostapchenko, A., & Schultz, H. (2007). Improving autonomous estimates of dem uncertainties by exploiting computer matching asymmetries. *8th Conference on Optical 3-D Measurement Techniques*. Zurich, Switzerland.
- Corrada-Emmanuel, A., & Schultz, H. (2008). *Autonomous estimates of horizontal decorrelation lengths for digital elevation models* (Technical Report). Computer Science Department, University of Massachusetts at Amherst.
- Donoho, D. L. (2006a). Compressed sensing. *IEEE Transactions on Information Theory*, 52, 1289–1306.
- Donoho, D. L. (2006b). For most large underdetermined systems of linear equations the minimal ℓ_1 -norm solution is also the sparsest solution. *Communications on Pure and Applied Mathematics*, 59, 797–829.
- Faugeras, O., Luong, Q.-T., & Papadopoulos, T. (2001). *The geometry of multiple images : the laws that govern the formation of multiple images of a scene and some of their applications*. Cambridge, Mass.: MIT Press.
- Hartley, R., & Zisserman, A. (2000). *Multiple view geometry in computer vision*. Cambridge, U.K. ; New York: Cambridge University Press.
- Pollefeys, M., Koch, R., & Gool, L. V. (1999). Self-calibration and metric reconstruction inspite of varying and unknown intrinsic camera parameters. *International Journal of Computer Vision*, 32, 7–25.
- Wakin, M., Laska, J., Duarte, M., Baro, D., Sarnovtham, S., and Kevin Kelly, D. T., & Baraniuk, R. (2006). An architecture for compressive imaging. *Proceedings of the International Conference on Image Processing – ICIP 2006*. Atlanta, GA.