'Photogrammetric Week '95', D. Fritsch & D. Hobbie, Eds., Wichmann Verlag, Heidelberg, 1995

The Effects of Image Compression on Automated DTM Generation

CRAIG ROBINSON, East Perth, BRUCE MONTGOMERY, Perth,
and
CLIVE FRASER, Melbourne

ABSTRACT

The effects of JPEG compression on automated DTM extraction via the approach of feature-based matching are investigated. JPEG lossy compression involves a truncation of higher spatial frequency data, a process which influences the accuracy of computation of grey-value gradients in the feature determination phases of image matching. Resulting accuracy effects on DTM heights obtained via the MATCH-T software system are investigated using a single stereomodel of 1:18,000 scale aerial photography at digital image resolutions of 15, 30, 45 and 60 µm. Heighting errors are computed for a range of compression ratios from 2:1 to about 40:1, illustrating that the impact of compression on DTM accuracy can be significant. It is concluded from this study that it is preferable to use the highest resolution available and to compress the image to the maximum extent allowed given the constraints of accuracy and image file size.

1. INTRODUCTION

Implementation of digital photogrammetric systems in a production environment is occurring at an accelerating rate. Indeed, the pace is such that operations such as semi-automated aerotriangulation,DTM extraction, orthophoto production and stereo digitizing are being routinely carried out in the absence of well-tested empirical models covering the effects of various aspects of digital photogrammetric processes upon the final mapping product. Traditional rules of thumb exist, for example sampling theory would suggest pixel sizes of 7-15 µm to match 60-30 lp/mm resolution on film (Jaakkola & Orava, 1994), orthophotos need a resolution of 3-8 pixels/mm at display scale (e.g. Leberl, 1994), and semi-automated DTM extraction from frame photography can yield accuracies to 0.1 ‰ h, where h is the flying height (Krzystek, 1991). Verifiable relations, however, are currently conspicuous by their absence.

Moreover, uniform agreement on vague rules of thumb has not yet been reached. Consider the question of image pixel size on DTM extraction for example. Implicit in the accuracy of 1/10,000 of flying height is a certain scanning resolution, which from basic geometric considerations should be in the order of 18 µm (assuming mensuration accuracy of 0.3 pixel). Results from controlled tests, however, do little more than to suggest that this specification is reasonable. Madani (1993) has reported that a 15 µm pixel size yielded DTM accuracies with the MATCH-T software package (Krzystek, 1991) which surpassed this rough measure by 20%; whereas 30 µm pixels yielded results which were poorer by 20% than the 'theoretical values'. The range reported by Krzystek & Wild (1992) is even wider, being from 40% better DTM accuracy than expected for 15 µm scanning, to 10-20% better for 30 µm pixels. This would suggest a matching precision of better than 0.3 pixels. Given the number of variables involved, ranging from image quality and content, through to the model for DTM approximation from an irregular object point field, these accuracy variations are not unexpected, though they are disconcerting to the photogrammetrist seeking to achieve an optimal design for a DTM generation project.

Another variable in the DTM accuracy equation is image compression, which is a viable means to partially alleviate the storage and manipulation problems associated with high resolution digital images of 23 x 23 cm aerial photography. Here too, vague rules of thumb seem to abound: "file reductions of 3 to 4 are common" (Boniface, 1994), "the compression method did not affect significant changes on the image geometry for ratios of 1:1 to 1:10" (Jaakkola & Orava, 1994), and "the range of useful compression ratios for JPEG is between about 3:1 and 20:1. For ratios of more than 20:1, the block

'Photogrammetric Week '95', D. Fritsch & D. Hobbie, Eds., Wichmann Verlag, Heidelberg, 1995

artefacts resulting from the division of the image into 8 by 8-pixel blocks become highly objectionable" (Villasenor et al, 1993).

Image compression, and most commonly JPEG still-image, lossy compression forms an integral part of today's high-end softcopy (digital) photogrammetric systems, where it is important for two further reasons beyond the obvious one regarding data storage (Toth, 1994). The first is that data access times and transfer rates are proportional to the actual data file size, and the second is that data size significantly influences processing time. The question of acceptable compression ratios is typically addressed in the context of visual assessment of the reconstructed image following lossy compression. While this may serve the needs of manual image mensuration, it does not necessarily provide insight into the impact of JPEG on automated image mensuration through centroiding or matching, be it feature- or area-based.

In this paper, the effects of JPEG compression on DTM generation via the MATCH-T software system are investigated. The scope of the investigation is modest, being confined to a single stereo model of 1:18,000 scale aerial photography, which has been scanned to yield a range of pixel sizes from 15 µm to 120 µm. Only resolutions of 15, 30, 45 and 60 µm are addressed in the present paper. The two higher resolutions (uncompressed) have been shown to yield DTM accuracies surpassing 0.1 ‰ h (Madani, 1993; Krzystek & Wild, 1992), whereas the two lower resolutions were included since they can be utilized in some instances for DTM generation in support of orthophoto production. The compression ratios considered ranged from 2:1 to about 40:1, and all processing was carried out on an Intergraph ImageStation Photogrammetric Workstation incorporating JPEG compression hardware and MATCH-T. With the investigation being limited to the single aim of assessing JPEG compression effects alone on DTM accuracy, and being confined to a single model, the realistic goal was to produce findings that would support practical guidelines rather than any definitive behavioural model linking compression and the accuracy of DTM extraction.

2. JPEG COMPRESSION AND MATCH-T

The JPEG still-image, lossy compression standard formulated by the ISO Joint Photographic Experts Group is currently the favoured compression algorithm for digital photogrammetric applications. Apart from being a well-known standard, JPEG has practical advantages in that it supports fast, 'on-the-fly' compression and decompression with the algorithms being capable of hardware implementation (e.g. Toth, 1994). Under JPEG the image to be compressed in subdividedinto 8 x 8 pixel blocks, with the image blocks being converted to the frequency domain via a discrete cosine transform. A quantization process then takes place in which the transformed coefficients are scaled and converted to integers. It is in this 'lossy' step only that the geometric and radiometric quality of the image is degraded. A lossless entropy encoding step follows quantization to complete the compression process (e.g. Aravind et al, 1993).

Within the quantization process a user-specified scaling factor, the q-factor, governs the degree of removal of spatial frequency data and therefore compression accuracy. Experience suggests that an acceptable range of q-factors would include values which yield compression ratios of between 3:1 and 20:1 (Villasenor et al, 1993; Aravind et al, 1993). This compression ratio range, however, is selected on the basis of minimising visible distortion, which comes in the form of block artefacts. These artefacts result from the 8 x 8 pixel blocking and they have the potential of compromising computer-based image processing tasks which operate on the reconstructed images. False edge effects complicate image analysis and segmentation, as well as edge detection (Villasenor et al, 1993).

Automated DTM extraction via MATCH-T involves two distinct parts: a pre-processing stage and the actual DTM generation step (Krzystek & Wild, 1992). It is in the first part, and especially in the formation of feature pyramids, that the impact of image compression will manifest itself. At each level of the image pyramid, the associated feature pyramid comprises image point features derived by the

'Photogrammetric Week '95', D. Fritsch & D. Hobbie, Eds., Wichmann Verlag, Heidelberg, 1995

Förstner operator (Förstner, 1986; 1993). The window selection and feature location steps of interest value computation involve the use of pixel intensity gradients, which are in turn influenced by the quantization process within JPEG compression. While these effects may be beyond the level of human visual perception, the truncation of higher spatial frequencies will certainly influence the computation of grey-value gradients and potentially degrade the geometric integrity of the feature location. The question is by how much. One instance where JPEG compression effects could be expected to be most pronounced is when the border of an 8 x 8 pixel block of a highly compressed image falls within the pixel window (5 x 3 for epipolar geometry) used for the interest operator computation. Within the reported investigation no attempt has been made to analyse compression effects on feature extraction and subsequent matching from a theoretical standpoint. Instead, a limited practical test was carried out in order to gain a quantitative 'feel' for the degree of DTM accuracy degradation accompanying the information-losing step of JPEG compression.

3. EXPERIMENTAL PROCEDURE

Figure 1: Ortho-rectified Image of Guilderton Test Site.

'Photogrammetric Week '95', D. Fritsch & D. Hobbie, Eds., Wichmann Verlag, Heidelberg, 1995

Within the workflow that produces a DTM from a stereo pair of absolutely oriented, epipolar-resampled aerial images, a number of error sources combine to influence the final precision of the height information determined. In seeking to investigate the influence of image compression alone on MATCH-T feature-based matching, it was necessary to design an experiment in which the errors in the photogrammetric orientation, resampling and DTM formation stages could be minimised. This was achieved by establishing a 'benchmark' DTM and then comparing subsequent DTM data obtained via JPEG compressed images to this 'ground truth' data, which was obtained from non-compressed data at the highest resolution level considered (15 μm). Figure 1 shows an ortho-rectified image of the test site selected for the investigation. The figure corresponds to the area covered by a single stereopair of 1:18,000 scale photography obtained with a wild RC10 camera of 152 mm focal length over the Guilderton area of Western Australia. The site comprised gently undulating terrain and sparse vegetation, the latter characteristic being chosen to minimise the density of ground surface 'disturbances' such as the vegetation canopy and buildings. Within the imagery there was a reasonable distribution of areas of high, medium and low spatial frequency. Diapositives were scanned on the Intergraph/Zeiss PS1 Photoscan system, contrast balanced, at a resolution of 15 μm. Hardware storage limitations at the time of the investigation precluded consideration of 7.5 μm scanning.
The general methodology for the practical testing was as follows:

a) A DTM was constructed from the uncompressed, 15 μm resolution epipolar-resampled imagery via MATCH-T. This DTM was held as the 'true' elevation data against which terrain heights obtained in subsequent determinations with compressed imagery were to be compared.
b) The same imagery was then compressed utilizing 22 q-factor values ranging from 10 to 799 (approximate compression ratios of 2:1 to 40:1).
c) DTMs were then derived for each of these compressions, with all model orientation and MATCH-T parameters being held fixed at the values corresponding to the benchmark determination. A 20 m grid interval was utilized in conjunction with a border delineating the nominal 3.1 km x 1.8 km area of interest. Each DTM comprised 13950 points.
d) The 15 μm uncompressed imagery was then sub-sampled, using an averaging algorithm for the pixel aggregation to produce images at 30, 45 and 60 μm resolutions, and the procedures of b) and c) were repeated. (This phase of the testing was also extended to resolutions of 75, 90, 105 and 120 μm; for a summary of these results, which are not reported here, refer to Robinson, 1994.)
e) Deviations in DTM height data for each resolution and degree of compression were then evaluated against the benchmark DTM, thus yielding both an error surface 'map' and an estimate of RMS heighting error. Recall that this RMS error is relative to the uncompressed, 15 μm resolution DTM and does not include components associated with the remaining photogrammetric processes.

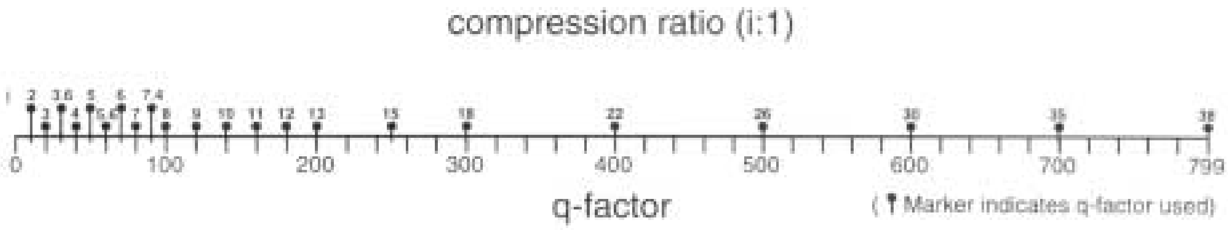

Figure 2: Approximate values of compression ratios for different q-factors.

Of interest to the photogrammetrist dealing with JPEG compressed images is the data storage savings associated with different values of the quantization scaling parameter. Unfortunately the compression ratios associated with different q-factor values are data dependent, the degree of compression being a function of resolution and the spatial frequency distribution within the image. For example, a q-factor

'Photogrammetric Week '95', D. Fritsch & D. Hobbie, Eds., Wichmann Verlag, Heidelberg, 1995

value of 30 leads to compression ratios of 3.8:1, 3.6:1 and 3.5:1 for one of the test images, for resolutions of 15, 30 and 45 µm, respectively. Corresponding figures for a q-factor of 100 are 8.3:1, 7.9:1 and 7.6:1. In order to give an indication of the compression ratios achieved, approximate values related to the 8-bit test imagery are given for different q-factors in Figure 2.

Based on the rule of thumb suggesting a DTM heighting precision of 0.1 ‰ h, the uncompressed 15 µm imagery should yield elevation data to an accuracy of about 0.27m. Within the absolute orientation phase, checkpoint residuals indicated that triangulation accuracy from manual stereo observations was well within this tolerance. One error source that did prove to be of greater magnitude than anticipated, however, was the impact of the epipolar resampling algorithm. The RMS height discrepancies between DTM heights generated from the uncompressed 15 µm resolution imagery, using bi-linear interpolation and cubic convolution, amounted to 0.13 m. For all subsequent MATCH-T runs at the different resolutions and quantization scale factor values, bi-linear interpolation was chosen.

4. RESULTS

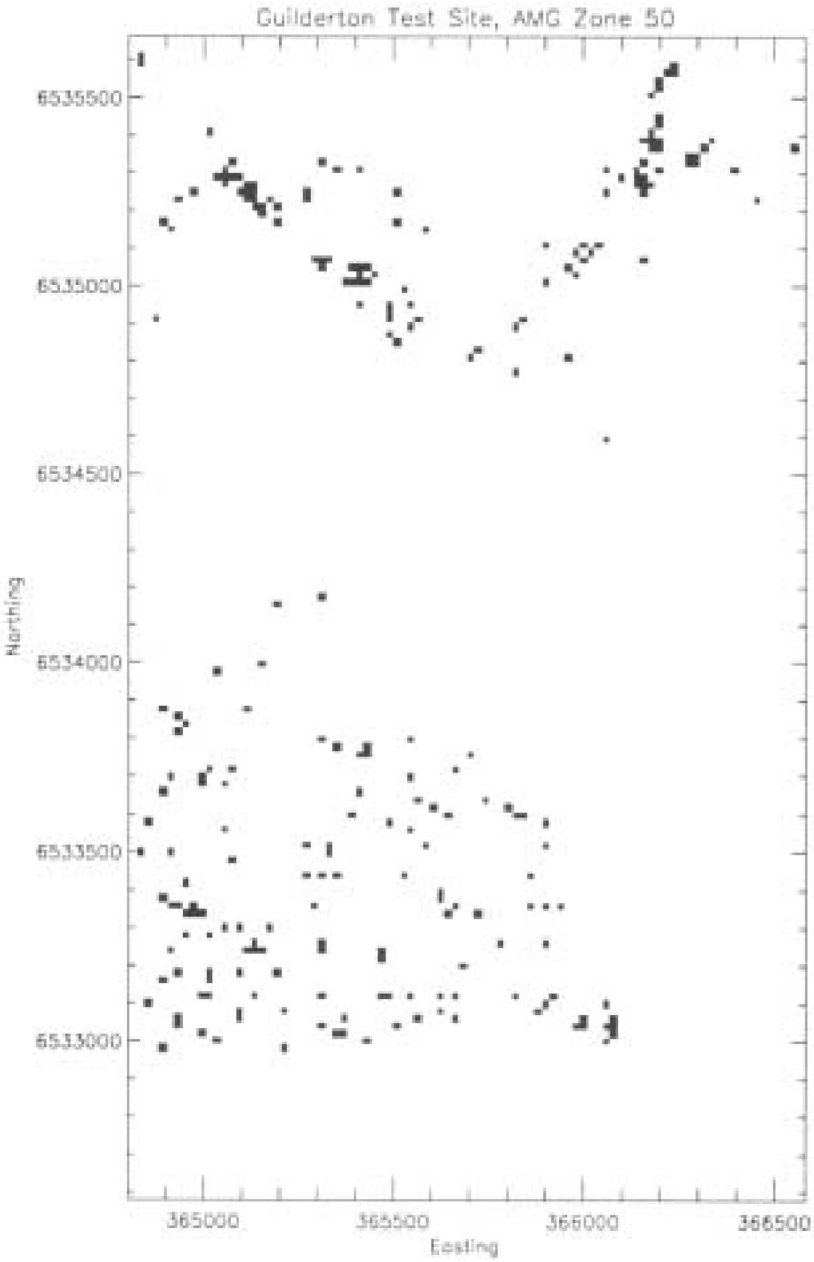

Figure 3: Spatial distribution of outliers for 15 µm imagery compressed at a q-factor of 10.

'Photogrammetric Week '95', D. Fritsch & D. Hobbie, Eds., Wichmann Verlag, Heidelberg, 1995

A total of 92 DTM computations were carried out for the four image resolution values. For each MATCH-T run height discrepancies against the benchmark 15 µm DTM were determined, and outliers identified. The number of outliers (> 3σ) present in each run varied from 1% (140) to 1.8% (250) of the total number of points, with the distribution in all cases being similar to that shown in Figure 3. Outliers tended to be concentrated in regions of high spatial frequency, for example in fallow fields and dry watercourses. Interestingly enough, the number of outliers did not vary significantly with either resolution or the degree of image compression.

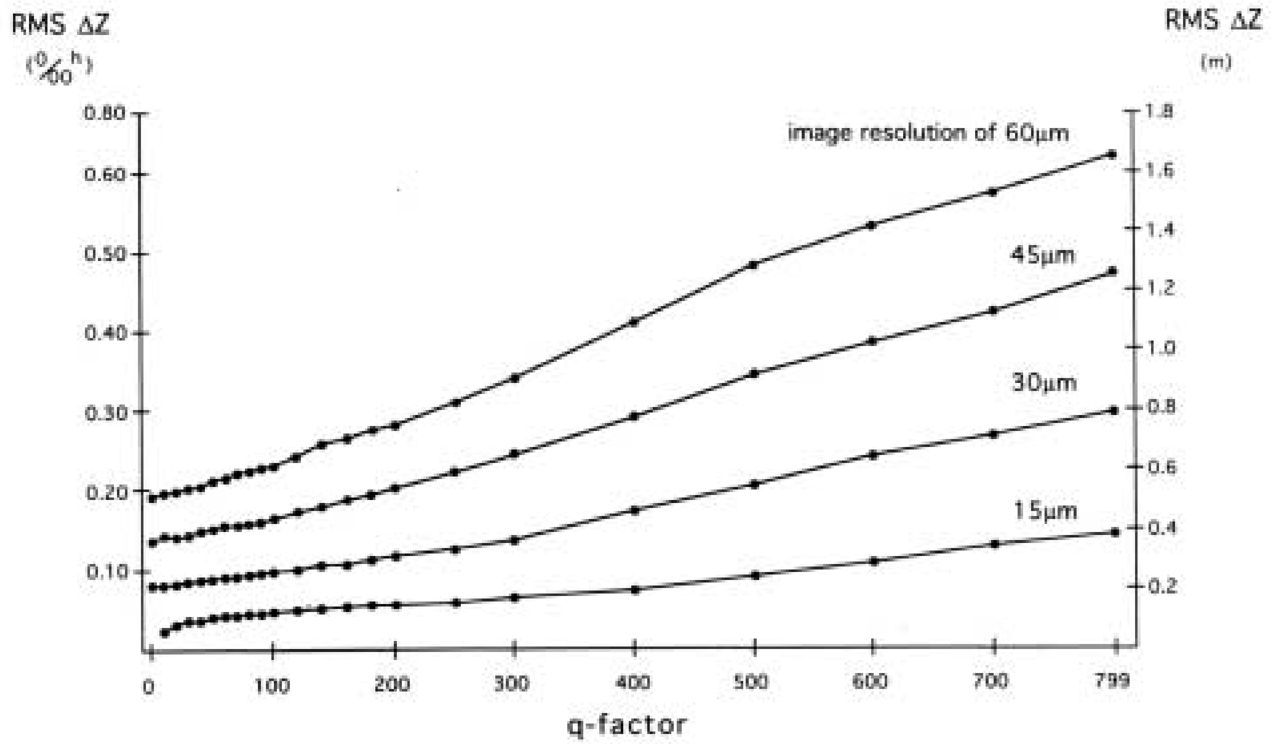

Figure 4: RMS heigthting errors relative to the 15 µm benchmark DTM for JPEG compressed images at four resolutions.

The principal results of the investigation are shown in Figure 4 which illustrates the variations in DTM accuracy (relative to the uncompressed 15 µm DTM) for the four image resolutions over the JPEG q-factor range of 10 to 799. A first inspection of the error curves suggests a near linear fall-off in accuracy with increasing q-factor values, the gradient being inversely proportional to image resolution. Prior to looking at the JPEG compression effects, however, it is interesting to note the linear change in the RMS heighting values for the uncompressed images at 30, 45 and 60 µm. Under the assumption of 1/3 pixel precision within the feature-based matching, the degradation in DTM accuracy for each 15 µm change in image resolution should be in the order of 0.15 m (1.6 x 18000 x 5 µm). And so it is between the three lower resolution image pairs. This analysis does not fully account for the 0.22 m RMS value between 15 µm and 30 µm resolutions, however.
The RMS error curves in Figure 4 provide an indication of the anticipated fall-off in relative heighting precision with increasing compression ratios and decreasing scanning resolution. To gain a practical measure of the absolute accuracy the RMS values must be combined with the 0.1 ‰ h precision associated with the benchmark DTM, thus yielding the standard errors σ_Z of DTM heights indicated in Figure 5. As an example, consider the case where a final DTM height accuracy of 0.2 ‰ h (0.5 m) is

'Photogrammetric Week '95', D. Fritsch & D. Hobbie, Eds., Wichmann Verlag, Heidelberg, 1995

required for the Guilderton test area. In this case, combinations of resolution and compression ratios meeting this tolerance are: all 15 µm determinations, 30 µm resolution with compression ratios of 20:1 or smaller, and 45 µm resolution at 6:1 or smaller. Each of these combinations yields a different image

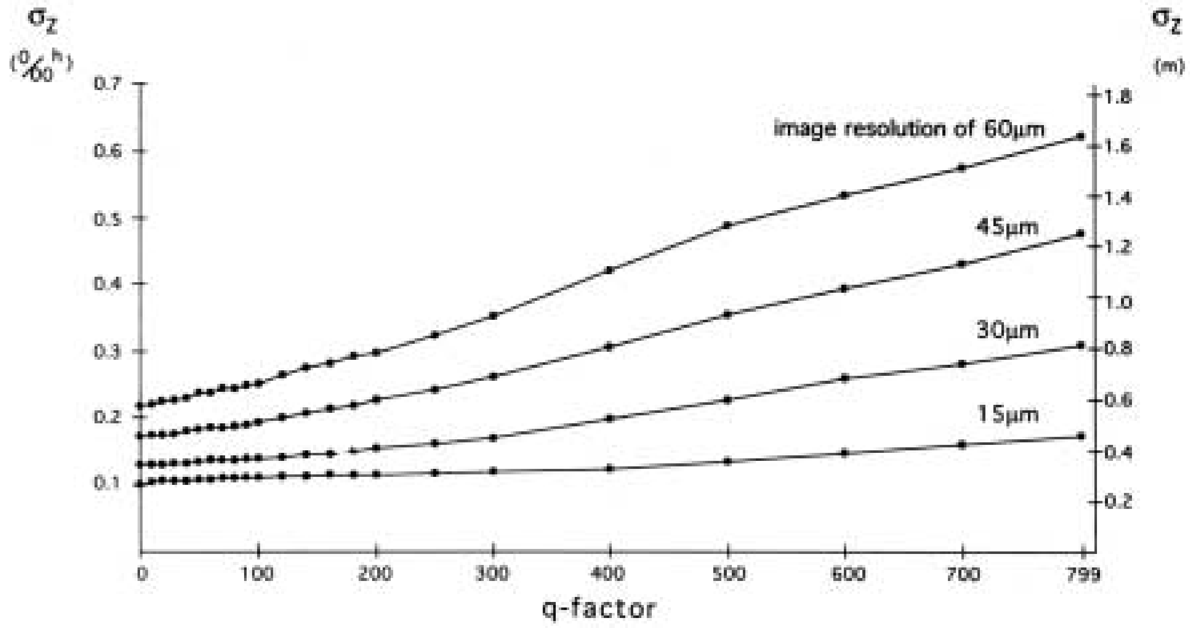

Figure 5: Absolute DTM heighting precesion, σ_z, inferred from relative heighting error (Figure 4) and assumed benchmark DTM accuracy of 0.1‰ h.

file size and thus the information can be used to optimize both data storage requirements and the associated computation times involved. The gradients of the curves in Figures 4 and 5 would suggest that for automated DTM extraction it is better to choose a highly compressed, high resolution image rather than a lower resolution image with a smaller compression ratio that yields the same data file size.

Computation times for the DTM extraction via MATCH-T are dependent on image resolution, and independent of the degree of JPEG compression. Thus, it would seem attractive from a productivity standpoint to always opt for the lowest resolution that will yield sufficient heighting accuracy. This is of course at odds with the previous recommendation, based on accuracy considerations, to choose higher image resolution.

5. CONCLUDING REMARKS

The results obtained in this investigation illustrate that JPEG compression has a significant influence on automated DTM extraction via a feature-based matching process such as that embodied in MATCH-T. The magnitude of the effect is dependent upon both image resolution and the degree of scaling of the quantization matrix in the JPEG compression operation. Relative DTM heighting accuracies fall off in a linear fashion with increasing q-factor values (i.e. higher compression ratios), the gradient of the accuracy degradation becoming steeper with decreasing image resolution. It is concluded that the truncation of the higher frequency data in the JPEG quantization, with its subsequent impact on image gradient values and therefore on interest operator computation and feature-based matching, is primarily responsible for the introduction of the heighting error.

The findings of this investigation support the conclusion that for automated DTM extraction it is preferable to use the highest image resolution available, and then compress with the maximum q-factor allowed given the constraints of accuracy and image file size. This recommendation is of course based on accuracy rather than productivity considerations which would favour lower image resolution.

Jaakkola and Orava (1994) have also concluded from their study of compression effects on least-squares matching that for small compression ratios of 10:1 or less, compression methods cause smaller image perturbations than increases in the pixel size. The standard error curves given in Figure 5 can prove useful in the selection of an optimal resolution/compression arrangement, though it should be stressed that this information relates to one particular stereomodel of 1:18,000 scale photography and may not be sufficiently representative of all cases. While it could be inferred from Figure 5 that a 7.5 mm resolution coupled with a compression ratio of up to about 20:1 might yield heighting accuracy surpassing the 0.1 ‰ h level, other factors can adversely influence automated DTM extraction from aerial photography scanned at such a high resolution.
As a final illustrative example of the effects of compression on DTM accuracy, consider the use of 15 mm resolution imagery with a q-factor value of 100 (compression ratio of approximately 8:1). Such a compression would generate a savings in disk storage of 88% for each image, from 240 mb to 29 mb. Yet, the change in the absolute DTM accuracy accompanying this degree of compression is only 0.03 m or 0.01 ‰ h. This figure is obtained by subtracting the assumed 0.27 m benchmark DTM accuracy from the combined influence of this value plus the RMS value of the relative error accompanying the JPEG compression (0.30 m). Depending upon accuracy specifications this increase might be seen as a small price to pay for such a savings in data storage requirements.

6. REFERENCES

Aravind, R., G.L. Cash, D.L. Duttweiler, H.-M. Hang, B.G. Haskell and A. Puri, (1993): Image and Video Coding Standards. AT & T Technical Journal, 72(1), pp. 67-88.

Boniface, P. (1994): State-of-the-Art in Softcopy Photogrammetry. Proc. Mapping and Remote Sensing Tools for the 21st Century, Washington, D.C., August 26-29, pp. 205-210.

Förstner, W. (1986): A Feature-Based Correspondence Algorithm for Image Matching. International Archives of Photogrammetry & Remote Sensing, 26(3/3), pp. 150-166.

Förstner, W. (1993): Feature Extraction in Digital Photogrammetry. Photogrammetric Record, 14(82), pp. 595-611.

Jaakkola, J. and E. Orava, (1994): The Effect of Pixel Size and Compression on Metric Quality of Digital Aerial Images. International Archives of Photogrammetry & Remote Sensing, Munich, 30(3/1), pp. 409-415.

Krzystek, P. (1991): Fully Automatic Measurement of Digital Elevation Models with MATCH-T. Proc. ISPRS Conference of Digital Photogrammetric Systems, Munich, September 3-6, pp. 203-214.

Krzystek, P. and D. Wild, (1992): Experimental Accuracy Analysis of Automatically Measured Digital Terrain Models. In Förstner/Ruwiedel (Eds.): Robust Computer Vision - Quality of Vision Algorithms, Wichmann Verlag, Karlsruhe, pp. 372-390.

Leberl, F.W. (1994): Practical Issues in Softcopy Photogrammetric Systems. Proc. Mapping and Remote Sensing Tools for the 21st Century, Washington, D.C., August 26-29, pp. 223-230.

Madani, M. (1993): A Preliminary Study on Accuracy of a Digital Photogrammetric Workstation. Presented paper, International Symp. on Operationalization of Remote Sensing, ITC Enschede, The Netherlands, April 19-23, 11p.

Robinson, C. (1994): The Effect of JPEG Image Compression on Digital Terrain Models Derived Using a Feature-Based Matching Algorithm. Unpublished MSc Thesis, Curtin University ofTechnology, Perth, 139p.

Toth, C. (1994): Image Compression in Photogrammetric Practise - An Overview. Proc. Mapping and Remote Sensing Tools for the 21st Century, Washington, D.C., August 26-29, pp. 70-79.

Villasenor, J., B. Belzer, and M.-J. Tsai, (1993): Compression of Digital Orthophotos. Final Contract report by UCLA Electrical Engineering Dept., USGS Western Mapping Center, Menlo Park, 19p.