ACCURACY ASSESSMENT OF AERIAL PHOTOGRAPHS ACQUIRED USING LIGHTER-THAN-AIR BLIMPS: LOW-COST TOOLS FOR MAPPING RIVER CORRIDORS

DAMIÀ VERICAT,* JAMES BRASINGTON, JOE WHEATON and MICHAELA COWIE

Centre for Catchment and Coastal Research, Institute of Geography and Earth Sciences, Aberystwyth University, Wales, Ceredigion SY23 3DB, UK

ABSTRACT

Monitoring river systems with repeat aerial photography is a powerful tool although the temporal resolution of surveys is rarely performed at anything better than annual time-scales. In recent years, a variety of low-cost aerial platforms for acquiring aerial photography have emerged. While these economical options may facilitate more frequent repeat surveys, the accuracy of imagery needs further consideration.

The accuracy of imagery obtained from a Lighter-Than-Air Blimp is investigated through two simple experiments. The first looks at the geospatial error of aerial photographs derived using five geometric transformation models, and the sensitivity of the photo registration quality to various ground control point (GCP) configurations and densities. At high GCP densities, higher order polynomial transformation models provide the highest quality registrations. However, at more modest GCP densities (i.e. 19–28 GCPs Ha^{-1}), simple aerotriangulation and 2nd order polynomial transformation models perform modestly, resulting in registration errors at standards equal to or better than obtained with conventional aerial photography (e.g. 0.5–1 m). The quality of image registration is highly dependent on the configuration of GCPs. In a second experiment, the practical utility of producing a mosaic of blimp acquired imagery is explored over a kilometre long braided reach. Even at relatively low GCPs densities (e.g. 9 GCPs Ha^{-1}), a mosaiced aerial of the entire reach can be produced of adequate quality to support bar-scale mapping of patch-scale features. The survey required less than a single day of field work and laboratory processing, and presents a cost-effective alternative to traditionally commissioned flights. Copyright © 2008 John Wiley & Sons, Ltd.

KEY WORDS: aerial photography; blimp; remote sensing; GCP; registration errors; river monitoring; repeat surveys

Received 29 April 2008; Revised 26 August 2008; Accepted 9 September 2008

INTRODUCTION

The last decade has witnessed significant advances in river science through the analysis of remotely-sensed imagery. This progress reflects important technological developments generated by; (a) new observation platforms (e.g. unmanned airborne vehicles, see Lejot *et al*., 2007); (b) new sensors (e.g. terrestrial laser scanning, see Hodge *et al*., 2008) and (c) new image processing algorithms (e.g. the HAB model of Fonstad and Marcus, 2005). These advances have dramatically extended the spatio-temporal resolution of remotely acquired data and now offer the prospect of system-scale datasets at very high spatial and spectral resolutions (Mertes, 2002).

The spatial coverage, high-resolution and historical availability of optical aerial photography is such that aerial photography remains central to research and applications in fluvial geomorphology (e.g. Surian and Cisotto, 2007) and stream ecology (e.g. Booth *et al*., 2007). Applications of optical aerial image analysis range from the characterization of in-stream parameters such as water depth, turbidity, substrate-type and algal concentration to floodplain-scale assessment of channel migration patterns and the systematic classification of fluvial and riparian habitats and community-level ecological mapping (see Marcus and Fonstad, 2008, for a comprehensive review). This broad range of uses has been facilitated by the availability of aerial photography in digital form, either by direct acquisition with scanning or CCD cameras, or through post-hoc digitization. In digital form, numerical image analysis can be used to extend traditional visual interpretation of hardcopy images and increase the scope of

---

*Correspondence to: Damià Vericat, Centre for Catchment and Coastal Research, Institute of Geography and Earth Sciences, Aberystwyth University, Wales, Ceredigion SY23 3DB, UK. E-mail: ddv@aber.ac.uk

information retrieval. For example, local image texture has been used to automate mapping of fluvial substrate size (e.g. Verdu *et al*., 2003; Carbonneau *et al*., 2005; Verdu *et al*., 2005,) and developments in softcopy stereo-photogrammetry has facilitated the production high-quality laboratory and field-scale DEMs of river systems (e.g. Brasington *et al*., 2003; Brasington and Smart, 2003).

Low-cost aerial platforms

Significant collections of airborne optical imagery are now readily available through government and commercial archives. However, the cost of acquiring repeat photography for bespoke applications and to facilitate routine monitoring (e.g. flood-scale mapping of channel migration or habitat evolution) is substantial, given the minimum deployment costs associated with commissioning an appropriately equipped aircraft. Particularly when high spatial resolution is required, low-altitude flights or use of hyper-resolution cameras can drive costs up. In the face of these logistical and financial constraints, researchers are increasingly turning to private low-cost platforms to acquire imagery. Stationary tethered balloons have been used for some time (e.g. Vetrella *et al*., 1977; Church *et al*., 1998), but there is a demand for mobile platforms allowing greater spatial coverage. Such mobile platforms include fixed-wing and helicopter UAVs (unmanned aerial vehicles), kites and lighter-than-air (e.g. helium) balloons or blimps. Very versatile drones (e.g. Lejot *et al*., 2007) and balloons and blimps exist (e.g. SWAMI with a payload of >300 kg, Chen and Vierling, 2006; Vierling *et al*., 2006), but these are not low-cost solutions. Of interest here are lightweight platforms that are readily controllable, easily deployed in remote locations and have low running costs making them ideal for customized monitoring projects. Such solutions can also be operated at very low altitudes to generate sub-centimetre resolution products capable of resolving grain-, cluster- and patch-scale morphologies (e.g. Church *et al*., 1998) well beyond that obtainable with even the most expensive digital aircraft-borne sensors.

Photography from lightweight aerial platforms, however, necessarily involves compromises due to the tight constraints on the operating payload that must accommodate optics, data storage/relay, actuators, plus any orientation and navigation system. While larger motorized UAVs can carry moderate payloads of 2–10 kg (e.g. 4 kg for the Pixie drone used by Lejot *et al*., 2007), small, easily portable and low-cost blimp-based photographic systems are typically constrained to operating loads of between 0.25–0.5 kg per 1 m^3 of lighter than air gas. This unforgiving design limit hinders the incorporation of standard positioning-orientation (i.e. coupled GPS and Inertial Navigation Systems or GPS-INS) and gyro-stabilization systems employed on conventional airborne platforms. As such, standard low-cost blimp-based photography is typically acquired using low-grade uncalibrated compact cameras. Such photography suffers from unknown exterior orientations (position and attitude) and pitch angles (i.e. oblique as opposed to truly vertical photography).

Image rectification

With such poorly constrained acquisition, the process of rectifying images to a common projection and coordinate system is critical for the conversion of blimp-based aerial photography to map-grade survey products. While true of all aerial surveys, the unknown exterior orientation and small footprint of low-altitude blimp photographs, exacerbates this standard image-positioning problem.

The process of co-registering photography to real-world coordinates can be achieved using a variety of transformation techniques, including aerotriangulation (sometimes termed Rubber Sheeting), polynomial georetification and orthorectification (see Wolf, 1984 or Lillisand and Kiefer, 2000 for an overview). However, given the typically low-relief of alluvial corridors and scarcity of quality elevation data, it is common practice to ignore the effects of elevation distortion.

Both aerotriangulation and polynomial georectification involve the transformation of image-referenced coordinates (row and column addresses of image pixels) to real world, or object-space coordinates (either x y Cartesian or easting/northing grid coordinates) through a set of ground control points (hereafter GCPs) clearly visible in the target image. GCPs are often deployed and surveyed in the field prior to image acquisition, but may be acquired after image acquisition provided there is sufficient recognizable geometry. Aerotriangulation and polynomial georetification differ, in that the former forces an exact match between image and object-space coordinates and then employs piecemeal triangular interpolation of pixels between three GCPs. In contrast,

polynomial georetification involves fitting a global linear, quadratic, cubic or higher order polynomial between GCP coordinate pairs, using a minimizer such as least-squares.

The accuracy of image co-registration is jointly dependent on: (i) camera orientation: in particular the camera attitude and degree of optical image distortion; (ii) the accuracy, density and distribution of GCPs; and (iii) the topographic complexity of the scene which controls the degree of image foreshortening and can only be adequately corrected by orthoretification. The significantly off-vertical (ultimately oblique) photographs acquired with low-grade compact cameras from blimp-platforms are likely to require a high-order transformation or precise aerotriangulation to rectify the high levels of distortion. The success of such extreme image transformation is, therefore, highly dependent on the quality of the GCP network available to constrain the solution. The logistical costs of deploying and recovering a high density of precision targets in the field can become significant, potentially offsetting the cost savings made by switching to lightweight aerial platforms. Thus, the identification of the optimal number and distribution of GCPs necessary to rectify low-altitude aerials and establish the accuracy of registered images is of paramount practical importance.

Objectives

With increasing numbers of researchers using amateur aerial photographic systems to generate low-cost image products, it is timely to evaluate the field- and laboratory methods necessary to generate photo-mosaics from blimps and provide some assessment of their reliability. In this paper, we describe two field experiments which aim to identify the magnitude and controls on geospatial errors in rectified blimp-based photography and propose simple methodological protocols for field acquisition and data processing. All the imagery described herein was acquired using a using a low-cost ($<$ UK £1000), $1.6 \, \text{m}^3$ helium blimp (described below), which has a maximum payload of $\sim$0.3 kg in normal flying conditions. Specifically, we explore:

1. The geospatial error of blimp-based aerial photographs derived using five geometric transformation models.
2. Sensitivity of transformation models to the density of ground-control points.
3. Sensitivity of transformation models to the spatial distribution of ground-control.
4. The development and evaluation of a field-to-laboratory protocol for mapping reach-scale fluvial features using blimp-based photography.

Below we report two field experiments designed to address the objectives listed above. The first of these (Experiment 1) involves an idealized case study over a flat, featureless surface (a sports pitch) with a dense network of high-precision GCPs. By systematically degrading the network of GCPs and using redundant points as check-data, this setup is used to evaluate the key controls (Objectives 1–3) on the rectification of a single image acquired with a low-cost blimp-based camera system. The second study (Experiment 2) comprises a more demanding field test, in which we acquire and register over 45 images covering a 1×0.4 km reach of a braided gravel-bed river. Using methods established in the idealized experiment, these images are georeferenced to create a very high resolution (5 cm) photomosiac of the study reach from which key features can be mapped within a GIS. Finally, we examine the reliability of derived map-data by comparing the pattern of predicted channel outlines with ground survey data using RTK GPS (Objective 4).

INSTRUMENTATION

The aerial platform

The term blimp is used here to describe any non-rigid airship that relies on high pressure inflation with a lightweight gas (typically helium) to provide buoyancy. The term was coined by the British Army in World War I, as a condensation of 'Balloon Limp', although it is also reputed to be onomatopoeic, reflecting the sound of tapping the inflated balloon. In these experiments we used a SkyHook Helikite© manufactured by Allsopp Helikites. This instrument comprises a $1.6 \, \text{m}^3$ helium filled envelope which exploits wind-generated lift through the incorporation of a single wing kite, as shown in Figure 1.

Figure 1. The SkyHook helikite. A $1.6\,m^3$ combined blimp-kite, with simple pan-tilt camera mount, secured by a ground-based tether. The blimp-kite is operated by a single person on the ground using a reel to control the height of the kite and walking to control its position. This figure is available in colour online at www.interscience.wiley.com/journal/rra

The kite not only provides lift, but also improves stability and extends the range of operating conditions. With appropriate air traffic permissions (necessary for flying heights exceeding 60 m in the UK), this platform can be launched in windspeeds between $0–13\,m\,s^{-1}$ and achieve a height of over 450 m. The helikite is manually controlled from the ground using a high tensile strength tether, which is easily released and rewound using a reel. Importantly, this versatile design provides a relatively high payload size, yet the blimp itself is easy to transport even to sites with difficult access. When not inflated, the blimp can be transported in a duffel bag, leaving the most difficult thing to transport being the lighter-than-air gas canister. Even when inflated, the platform is small enough $(0.9 \times 1.2 \times 1.7\,m)$ to be transported in a back of an estate car or short wheel-based van, and can typically maintain its inflated state for two to three successive field days with only minor topping up. Even so, the use of blimps in river corridors may be limited by external factors such as power lines or bridges. A reconnaissance survey of the study reach is therefore necessary prior to deployment to identify objective aerial obstacles such as power/communication lines and overhanging trees which may present snag hazards for the tether cable, and any additional ground obstacles that may prevent the mobility of the handler on the floodplain.

The camera

While remote-control (zoom, shutter-release etc.) of camera systems is possible using radio or wireless telemetry (Vierling *et al.*, 2006), a more cost-effective and lighter weight solution to acquiring photography from an aerial platform is to use continuous shooting or timed shutter-release to acquire a massively redundant dataset. Here we used a *Ricoh Caplio R5* (7.2 MP) digital camera, secured and orientated using a pan-tilt adapter attached to the base of the helikite. Manual trial and error was used to calibrate the downward look-angle required to achieve near vertical photography and compensate for the attitude of the helikite. The camera was set in auto-exposure mode, and has an in-built intervalometer, set to acquire an image every 5 sec. Using a 2 Gb internal SD storage card, over 30 min of photography can be acquired at the maximum image resolution (2084×1536 pixels) stored using lossless JPEG image compression (> 360 images). Manual visual filtering is then used to select images with optimal exposure and limited distortion; an important first step therefore, in minimizing the rectification problem (typically takes 15–30 min).

Pre-flight camera calibration

As with standard airborne photography, the trade-off between image object space resolution (the ground dimensions of an image pixel) and image ground coverage (the photographic footprint) is dictated by flying height. While the desirable image characteristics are clearly project specific, pre-flight planning is necessary to achieve optimal results. The relationship between flying height and image footprint is controlled by the camera optics and can be calculated by measuring the pixel separation distance for image targets with known object-space coordinates over a set of ranges (the equivalent of flying height). This calibration was established for a focal length of 28 mm using two photo-targets separated by 6.5 m in the horizontal over ranges of 7–35 m. Target positions and distances were all measured using a Leica 1200 total station. The resulting calibration used in these experiments is described in Equations 1 and 2, and assumes square pixels.

$$W_f = 1.14 \times R - 0.11 \tag{1}$$

where W_f is calibrated image horizontal width and R is the range (both, m); and

$$P_x = 0.00037 \times R - 0.00004 \tag{2}$$

where P_x is the pixel dimension (assumed to square) in metres.

This calibration provides a guide to the necessary flying height to achieve project specifications. For example, truly vertical photography acquired at a height of 200 m would delivery imagery with a ground footprint of 230×170 m and object space pixel dimensions of 7 cm. Note, that it is unlikely that the CCD array of the camera is truly square, so the vertical-horizontal dimensions should be taken only as an approximate guide.

Ground-control targets

GCPs for both sets of field experiments were deployed prior to photographic acquisition. These targets comprised 1 m squares of 4 mm thick black plastic sheeting, with photogrammetric markings superimposed in yellow reflective tape and labelled with a unique identifier. The position of target centroids was established using a Leica 1200 total station (during the Experiment 1) and Trimble R8 RTK GPS (during the Experiment 2). The 3d accuracy of GCPs coordinates thus varies between 0.003 m for Experiment 1 and 0.02 m for Experiment 2.

IMAGE PROCESSING

Image co-registration was undertaken in ERDAS IMAGINE 9.1. This software provides a visual environment to confirm pair-wise ties between the image pixel coordinates of GCPs and their corresponding object space position. A variety of transformation models can be fitted using this software. In this paper, we consider only polynomial models fitted by least squares minimization and exact aerotriangulation with piecemeal interpolation between triangles defined by three GCPs. Following successful transformation, an output image can be calibrated by resampling the transformed image pixels to a uniform spatial resolution which should neither under- or over-sample the original image. Various options for resampling pixel attributes (RGB brightness numbers) are available including bilinear, cubic or nearest neighbour interpolation. The choice of model depends on the image application. For example, nearest neighbour interpolation preserves the spectral integrity of image pixels, but can result in jagged between boundaries, whereas cubic interpolation smoothes boundaries but averages spectral information between pixels.

The reliability of image transformations can be measured using the root mean square error (hereafter RMSE) derived for the set of paired coordinates used to fit the model:

$$\text{RMSE} = \sqrt{\frac{1}{n}\sum_n (\hat{x} - x_i)^2 + (\hat{y} - y_i)^2} \tag{3}$$

where n is the number of GCP pairs, $\hat{x}$ and $\hat{y}$ are the transformed (predicted) object space coordinates of each GCP and x_i and y_i are its observed (sampled, assumed true) object space positions. Where available, independent GCPs

not used in fitting the transformation can then be used as check-data (hereafter, Check Points or ChPs) and an unbiased RMSE term derived to validate the resulting transformation.

EXPERIMENT 1: MODELS AND GCP NETWORKS

Data acquisition and experimental design

42 GCPs were deployed over a flat area of homogenous grass cover, measuring 70×60 m with an average spacing interval (roughly paced) of 10 m in x and y. Precise coordinates of each GCP were acquired in a local coordinate system using total station survey. The helikite-camera system was deployed in a gentle moderate breeze ($\sim$3–4 m s^{-1}), released to an approximate flying height of 150 m and set to record images at 5 sec intervals with the focal length of 28 mm. The instrument was manually guided to acquire photography over the target area. On retrieval, the optimal image for later processing, was selected based on the maximum number of visible GCPs; 23 targets covering a ground area of 0.36 ha (Figure 2a).
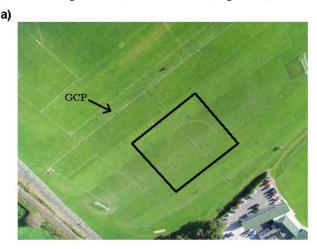
A simple experimental design was developed to explore the controls on image co-registration using five commonly employed transformation models; linear, quadratic, cubic and quartic polynomial and aerotriangulation (Figure 2b). First, the influence of GCP density on transformation model performance was assessed by systematically degrading the GCP network. In this experiment redundant GCPs were used as check points to calculate an unbiased RMSE error to describe the model fit. This results in a matrix of different image transformations using between 5–21 GCPs with between 2–18 check points for the five different models. Note that the minimum number of GCPs required to fit the polynomial models varies 3–7 for linear to quartic respectively.

Second, using the optimal model, the sensitivity of this transformation to the spatial distribution of GCPs within the scene was explored. This was analysed by segregating the GCPs into four spatial patterns: (i) uniform random spatial sample, (ii) biased to the image perimeter; (iii) biased to the image centre and (iv) distributed across the image diagonal. In each case, model errors were quantified using a network of nine uniformly distributed check points.

Results: 1. Sensitivity to GCP density

A summary of model results for the matrix of transformation scenarios is presented in Table I and shown graphically for each of the transformation models in Figure 3.

The summary of results provided in Figure 3 demonstrates the superior performance of high order, in particular the 2nd, 3rd and 4th order, polynomial models. In these cases, ChP errors (Figure 3a) are steady across a range of GCP densities and registration (i.e. GCP errors (Figure 3b)) errors are close to the optimal resampled pixel dimensions
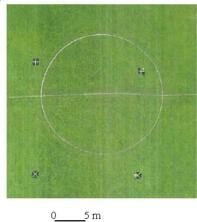

Figure 2. (a) Optimal image selected based on the maximum number of visible GCPs. A total of 23 GCPs cover a ground area of 0.36 ha. Inset black box depicts the area shown in b. (b) View of a selected part of the image after applying the aerotriangulation transformation model. This figure is available in colour online at www.interscience.wiley.com/journal/rra

Table I. Matrix of transformation scenarios designed to evaluate the sensitivity of transformation models to GCP density and registration errors for all transformation models

Scenario	Number of GCPs[1]	Density of GCPs (GCPs Ha^{-1})	Number of ChPs[2]	Density of ChPs (GCPs Ha^{-1})	Aerotriantgulation RMSE (m)[3]		Polynomial 1st order RMSE (m)		Polynomial 2nd order RMSE (m)		Polynomial 3rd order RMSE (m)		Polynomial 4th order RMSE (m)	
					GCPs[4]	ChPs	GCPs	ChPs	GCPs	ChPs	GCPs	ChPs	GCPs	ChPs
1	21	58	2	6	~0.00	1.08	4.78	1.18	0.51	0.53	0.10	0.09	1.37	0.07
2	20	56	3	8	~0.00	1.27	4.83	2.08	0.52	0.48	0.10	0.09	0.12	0.07
3	19	53	4	11	~0.00	1.12	4.91	2.19	0.53	0.50	0.10	0.08	0.48	0.07
4	18	50	5	14	~0.00	1.08	4.92	2.47	0.50	0.53	0.09	0.10	2.35	0.07
5	17	47	6	17	~0.00	0.89	4.98	2.71	0.52	0.49	0.09	0.11	1.00	0.07
6	16	44	7	19	~0.00	0.91	5.06	2.52	0.52	0.45	0.09	0.15	0.19	0.05
7	15	42	8	22	~0.00	1.08	5.04	2.95	0.49	0.51	0.09	0.12	0.17[5]	0.14[5]
8	14	39	9	25	~0.00	1.18	5.19	2.85	0.50	0.69	0.08	0.13	0.00[5]	0.00[5]
9	13	36	10	28	~0.00	1.26	5.33	2.91	0.47	0.70	0.07	0.14	[6]	[6]
10	12	33	11	31	~0.00	1.21	5.31	3.34	0.48	0.70	0.07	0.23	[6]	[6]
11	11	31	12	33	~0.00	0.97	4.72	4.72	0.45	0.45	0.02[5]	0.02[5]	[6]	[6]
12	10	28	13	36	~0.00	0.91	5.13	3.19	0.54	0.58	0.00[5]	0.29[5]	[6]	[6]
13	9	25	16	44	~0.00	1.30	5.27	3.27	0.54	0.61	[6]	[6]	[6]	[6]
14	8	22	15	42	~0.00	2.01	5.41	3.50	0.50	1.12	[6]	[6]	[6]	[6]
15	7	19	16	44	~0.00	2.71	5.37	3.64	0.49	1.25	[6]	[6]	[6]	[6]
16	6	17	17	47	~0.00	2.06	5.53	3.54	0.00[5]	1.91[5]	[6]	[6]	[6]	[6]
17	5	14	18	50	~0.00	4.57	5.47	4.15	[6]	[6]	[6]	[6]	[6]	[6]

[1]GCPs, Ground Control Points used for geometrically transformation.
[2]ChPs, Check Points used to test errors along the geometrically corrected photograph.
[3]RMSE, Root Mean Square Error (see methods for calculation).
[4]GCPs RMSE is almost zero (see text for more discussion).
[5]Results clearly affected by the small number of GCPs (see text for discussion).
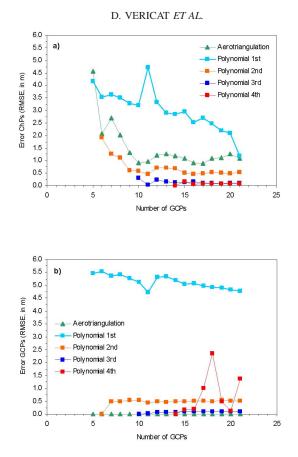[6]Not minimum GCPs available to apply the specified model transformation.

Figure 3. Number of GCPs versus the RMSEs of the ChP (a) and GCP (b) for different transformation models; (i) Aerotriagulation; (ii) First order (1st); (iii) Second order (2nd); (iv) Third order (3rd); (v) Fourth order (4th). This figure is available in colour online at www. interscience.wiley.com/journal/rra

(0.05 m). By comparison, the simple linear transformation model fails to provide a reliable, robust transformation of the image, with ChP and GCP RMSE continuing to fall with increasing GCPs. This model ultimately fails to achieve registration errors of less than 1 m, even with unrealistically high numbers of tie points. The best results overall, in terms of the ChP errors, are found for the 4th order model. However, unlike the 2nd and 3rd order models, this transformation also yields an irregular pattern of high GCP errors, suggesting a spatially unreliable pattern of transformation, varying strongly as new information is added or substracted. In contrast the lower order, non-linear models have commensurate GCP and ChP errors, albeit subject to a threshold minimum number of GCPs (e.g. 10 GCPs for both models).

With conventional aerial photography, acquired with calibrated optics and close to vertical from a stabilized airborne platform, standard linear (affine) or weakly non-linear transformation is usually sufficient to rectify imagery to an object-space coordinate system (see for example, Hughes *et al.*, 2006). In these cases, the transformed image coordinates remain parallel or close to parallel for 2nd order models, resulting in uniform or gently stretched imagery. The greater success of the higher order transformations found in this experiment reflects the need for highly non-parallel and curved image manipulation to adequately register the severely distorted (i.e. oblique) blimp photography. However, while such high order transformations are able to correct for such extreme distortion, marginal areas of the transformation poorly supported by the GCP network are liable to highly non-behavioural image skew; a pattern evident in the highly variable GCP errors evident in the quartic model. The results also highlight the need for a very dense GCP network to rectify imagery at levels of accuracy commensurate with the pixel size resolution. For example, optimal results obtained with the 3rd and 4th order models (typical ChP errors of 0.07–0.015 m) required over 15 GCPs, over the small image area, equating to some 45 GCPs Ha^{-1}

(Figure 3). Clearly, such densities are likely to be practically difficult to achieve for all but the most spatially restricted of projects.

Figure 3 also reveals the distinction between the exact co-registration provided by aerotriangulation (~zero GCP errors) and the least squares fitting used in the polynomial transformations. While this exact model yields negligible GCP errors, the ChP errors stabilize at a relatively high threshold error of ~1 m when applied with more than 9 GCPs (Figure 3). This comparatively high threshold error reflects the methodological dependence on piecewise linear interpolation between fixed GCPs. Unless GCPs are spaced at impractically high spatial densities, the linear interpolation cannot compete with the curved, non-parallel transformations applied to such severely distorted imagery.

Despite this, the overall results are encouraging with both aerotriangulation and the relatively parsimonious requirements of 2nd order transformation models shown to be capable of yielding registration errors at standards equal to or better than obtained with conventional aerial photography (e.g. 1 m) with a comparatively light GPC requirement of 19–28 GCPs Ha^{-1} (Table I and Figure 2).

Results: 2. Sensitivity to the spatial distribution of GCPs

The sensitivity of the transformation errors to the spatial arrangement of GCPs was analysed using data from Scenario 8 (see Table I) and applying the optimal 3rd order model. This scenario incorporates 14 GCPs and 9 ChPs and thus provides both the necessary the high minimum GCP demands of this high order model as well as sufficient independent check data to validate the results. Four spatial patterns of GCPs were analysed: (i) uniform random spatial sample; (ii) biased to the image perimeter; (iii) biased to the image centre and (iv) distributed across the image diagonal; all shown schematically in Table III. For all cases, ChPs were distributed uniformly across the image.

Table II. Individual GCP and ChP errors for the four spatial control network models. All results obtained using 14 GCPs and 9 ChPs and a 3rd order polynomial transformation

Position of the GCPs[1]							
Randomly[2] RMSE (m)[3]		Perimeter RMSE (m)		Centre RMSE (m)		Diagonal RMSE (m)	
GCPs	ChPs[4]	GCPs	ChPs	GCPs	ChPs	GCPs	ChPs
0.02	0.00	0.00	0.05	0.01	0.02	0.00	0.02
0.03	0.06	0.01	0.11	0.01	0.06	0.02	0.11
0.04	0.07	0.02	0.13	0.02	0.07	0.04	0.16
0.05	0.10	0.03	0.16	0.02	0.13	0.05	0.23
0.06	0.13	0.03	0.19	0.02	0.26	0.05	0.24
0.06	0.16	0.05	0.22	0.02	0.71	0.05	0.32
0.07	0.17	0.05	0.28	0.03	0.73	0.05	0.40
0.08	0.17	0.06	0.33	0.03	0.90	0.06	0.46
0.08	0.19	0.06	0.33	0.03	1.01	0.06	1.65
0.08	—[5]	0.06	—	0.04	—	0.06	—
0.10	—	0.07	—	0.04	—	0.06	—
0.11	—	0.08	—	0.04	—	0.07	—
0.12	—	0.08	—	0.05	—	0.07	—
0.13	—	0.12	—	0.08	—	0.08	—
Total RMSE (m)		Total RMSE (m)		Total RMSE (m)		Total RMSE (m)	
0.08	0.13	0.06	0.22	0.04	0.57	0.05	0.61

[1]GCPs: Ground Control Points used for geometrically transformation.
[2]Representative sketches where the square and the dots simulate the photograph and the GCPs respectively.
[3]RMSE: Root Mean Square Error (see methods for calculation).
[4]ChPs: Check Points used to test errors along the geometrically corrected photograph.
[5]9 ChPs have been used to perform this analysis, thus hyphened fields are not applicable.

The resulting spatial errors for the 14 GCPs and 9 ChPs (in rank order) are shown individually for each spatial scenario in Table II. This reveals the range of GCP RMSEs to be at a mimimum when tie-points are spatially concentrated, either in the image centre or along the diagonal. These patterns give rise to low total GCP errors, 0.04 and 0.05 m respectively and notably in the range of the output pixel cell-size. However, while the GCP errors are minimized for these scenarios, the independent check data reveal a counter pattern, with high total RMSE, reaching a maximum of 1.65 m for the diagonal scenario. In contrast, the random and perimeter spatial models reveal a more consistent pattern of more moderate GCP and ChP errors.

These results reflect the spatial bias that results when least squares minimization is used to fit the transformation model. When GCPs are spatially concentrated, an optimal local interpolation results, which while well constrained for the area within the GCP network, can yield a model with significant distortion in areas far from control. The use of a high order polynomial model exacerbates this problem and gives rise to severe unconstrained image curvature and high errors around the image perimeter.

Following the approach used by Hughes *et al.* (2006), the individual GCP and ChP errors can also be presented in the form of cumulative error probability plots. These allow the user to extract the probabilistic registration error at any given confidence interval and thus quantify the likely 'spatial risk'. This has particular value in establishing spatial error buffers or envelopes when analysing locational data derived from registered imagery. Figure 4 shows cumulative error probability plots for the four spatial models, quantified separately for the GCPs and ChPs.

This shows that the 0.5 probability GCP error rises from 0.03 to 0.08 m for the central to random arrangement respectively, while the 0.5 probability ChP error varies between 0.15 (random) and 0.73 m (centre distribution). The equivalent 0.9 probability error results show ChP errors increase up towards 1.1 m for the diagonal arrangement of GCPs. While only four spatial arrangements have been explored here, this simple statistical decomposition of the results provide an insight in the wide range of possible geospatial errors that can arise from inappropriate GCP coverages.

Discussion

The results from Experiment 1 suggest that co-registering blimp-based imagery to a high standard (horizontal accurcies of ±10 cm) requires both a high density GCP network and higher order models. Unlike conventional
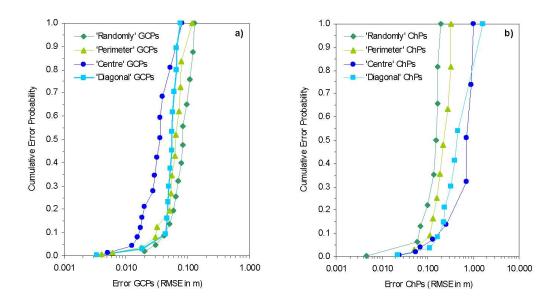

Figure 4. Cumulative error probability plots for the four spatial GCP patterns: (a) GCP errors; (b) ChP errors. This figure is available in colour online at www.interscience.wiley.com/journal/rra

aerial products, which can be registered typically with linear or low order models, the poorly constrained camera orientation resulting from the unstable blimp platform (Figure 2a), produces significant image warping and curvature to rectify the image to a planar coordinate system. With exceptionally high GCP densities (e.g. circa 40 GCPs Ha^{-1}) and an evenly distributed network of tie-points, higher (3rd and 4th) order models can produce relatively high quality imagery (i.e. co-registration errors commensurate with the obtainable pixel resolution: circa 10^{-2} m resolution). However, these demands will frequently exceed the logistical limits of most survey work as for example, a 1×1 km survey area would require nearly 4000 GCPs.

Using coarser GCP networks and less demanding aerotriangulation and 2nd order polynomial models yields absolute errors of 0.5–1 m (significantly larger than the obtainable pixel resolution). While errors of such magnitude may present challenges for applications like grain size extraction from imagery (e.g. Carbonneau, 2005, Verdu *et al.*, 2005) or construction of digital elevation models from stereo pairs, they are perfectly suitable for providing relatively high-resolution imagery for mapping and contextual purposes. This suggests that intermediate and larger-scale geometries (e.g. bar and channel scale in the case of fluvial applications) could be relatively well identified using this low-cost approach. The application of such a technique across a reasonable length reach will be explored in the second experiment in more detail.

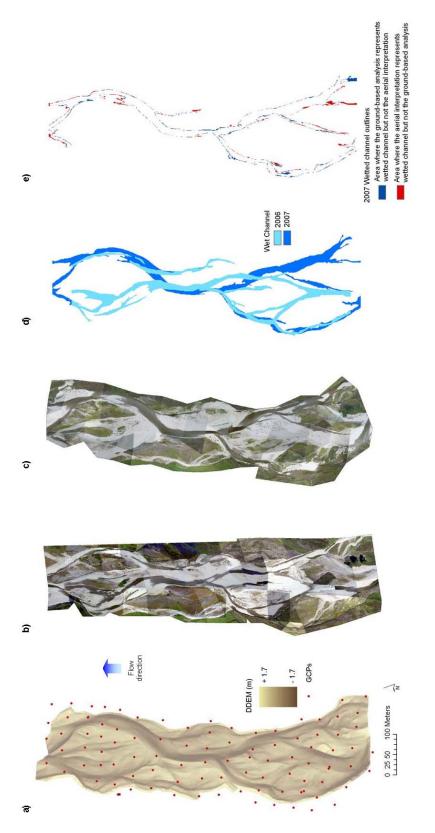## EXPERIMENT 2: FIELD SCALE AERIAL MAPPING

Geo-registered blimp surveys have a range of application in fluvial and wider geomorphological investigations, such as mapping facies or to monitoring rates and patterns of channel movement after formative events. Here we explore the reliability of blimp-image based interpretations of the pattern of channel kinematics between 2005 and 2007 on a $\sim 1 \times 0.3$ km reach of the braided River Feshie in Northern Scotland UK (Figure 5a). Changes in channel pattern of this river have been studied over different time scales by Werritty and Ferguson (1980) by means of the analysis of historical aerial photographs. Here we outline a simple protocol for mapping a reach of this river and evaluate the locational reliability of digitized map products (in this case, channel outlines) through comparison with ground-based survey using high precision RTK GPS.

Data acquisition

The scale of the study reach exceeds the footprint of a single blimp image flown in logistically feasible limits (e.g. up to e.g. 250 m flying height to yield e.g. 5–10 cm resolution imagery). Thus, acquisition over a larger area requires the systematic mosiacing of individually co-registered images. In order to expedite this process, 110 black GCPs similar to those used in Experiment 1 were deployed across the area of interested, at a roughly 50 m grid spacing interval (Figure 5a). These were then surveyed into a local planar coordinate system using RTK GPS (mean point precision ~ 0.02 m). The blimp was launched in moderate wind speeds in diffuse light conditions, with the camera orientated by trial and error to acquire near vertical photography, accounting for platform tilt in high wind. As in Experiment 1, the timed shutter release function on the Ricoh Capilio camera, was set to acquire an image every 5 sec, each stored as a lossless JPG image on a 2 Gb data card. Image acquisition over the study area was achieved by manually guiding the blimp along three longitudinal, overlapping flight lines. With the blimp set to an altitude of approximately 200 m, the object-space image footprint was in the region of 230×170 m, subject to wind variation. This flight plan ensured significant overlap between consecutive images, so that of the approximately 2000 images acquired in each annual survey (taking no more than 3 h of flying time), there was significant redundancy in the dataset to reduce this to 45 optimally exposed and near vertical images covering the study area.

Image processing

The resulting pattern of GCPs gave rise to a mean target density of approximately 9 GCPs Ha^{-1} and typically less than 6 GCPs per randomly orientated image (Figure 5a). This was too low to facilitate the application of more than 1st order polynomial models, so co-registration was attempted using aerotriangulation with piecewise interpolation as described above. Additionally, the sparse control network precludes the option to retain any GCPs from the transformation solution, so the only independent validation of co-registration process comes in the form of

Figure 5. (a) Detrended Digital Elevation Model (dDEM) with GCPS from the study reach in the River Feshie in 2007. Co-registered and mosaiced blimp imagery of a 1 km study reach of the River Feshie in (b) 2005 and (c) 2007. (d) Digitized wetted channel outlines for the mosaiced blimp imagery in 2005 and 2007. (e) Spatial intersection of the ground-based (RTK GPS survey) and aerial interpretation (digitized) channel outlines. This figure is available in colour online at www.interscience.wiley.com/journal/rra

comparing channel outlines digitized from the imagery with those obtained from ground-based survey (i.e. RTK GPS).

Individual images were co-registered and subsequently mosaiced using ERDAS IMAGINE 9.1. This software enables the user to select the order of images superimposed in an image mosaic, and this was selected on the basis on minimising exposure differences between scenes where possible. Finally, wetted channel outlines were digistized from the mosaiced image using manual image interpretation and ArcMap.

Results

Mosiacked images derived from the blimp surveys in 2005 and 2007 together with the pattern of digitized wetted channels are shown in Figure 5b, c and d. Despite the large number of images required to complete the mosaic, the final map product is of high quality (Figure 6), and required less than a single day of laboratory processing, again keeping costs low. The digitized wetted channel outlines reveal a significant channel avulsion to have occurred between 2005 and 2007, in the lower third of the study reach, with the dominant anabranch switching from true left to true right (Figure 5d).

Errors attributed to the image registration can be observed by comparing the ground-based (i.e. outline obtained by RTK GPS survey in the field) and the aerial interpretation (i.e. outline obtained digitizing) wetted channel outlines. Errors in the outline obtained digitizing vary throughout the image, and depend (in part) on the proximity to GCPs as might be expected for this transformation. However, in line with the results of experiment, errors are rarely found to exceed 1 m. Although wetted channel outline obtained with RTK GPS survey in the field may be considered as the 'true outline', errors associated with the precision to connect RTK GPS measurements may be on the order of 25–50 cm.

The quality of the predictive fit (aerial interpretation to ground-based outlines) can measured using a simple error metric, the F statistic, frequently used in the assessment of hydraulic model performance (e.g. McMillan and Brasington, 2008). This quantifies the spatial intersection of the observed and predicted channel outlines and yields a maximum score of 100% for error free predictions:

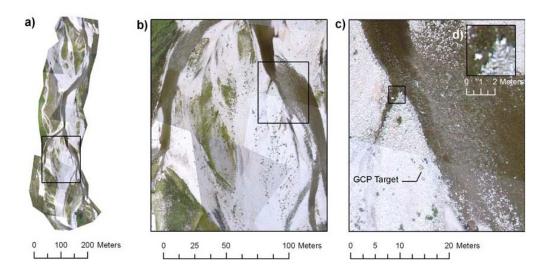$$F = \frac{A_{\text{obs}} \cap A_{\text{pred}}}{A_{\text{obs}} \cup A_{\text{pred}}} \tag{4}$$

Figure 6. Example of typical mosaiced imagery (registered applying aerotriangulation fit) at various zoom levels from 2007 on Feshie. (a) Overview of reach with inset black box depicting area shown in b. (b) Close up of sub-reach showing minor mismatch alignments between individually mosaiced images, with inset black box depicting area shown in c. Note that bar-scale morphology, waters edge, facies and vegetation can all be mapped reasonably at this scale. (c) View of bar-scale feature, with inset black box depicting area shown in d. Note the warping at the edge of the photo mosaic boundary across bottom 1/3 of image. (d) View at patch-scale showing lower limit of 5 cm resolution imagery. This figure is available in colour online at www.interscience.wiley.com/journal/rra

where A_{obs} is the observed wetted area measured by RTK GPS and A_{pred} is the predicted wetted area digitized from the aerial imagery.

Aerial interpretation and ground-based wetted channel outlines in 2007 have been rasterized (1 m grid-size) in order to be analysed using the Spatial Analysis extension of ArcMap. The raster file calculated by means of the aerial interpretation yields a wetted channel area of $20\,832\,m^2$ while in the case of the ground-based raster the area is $20\,907\,m^2$. The intersection between both raster files gives the possibility to calculate (i) the area (m^2) where both files represent wetted channel; (ii) the area (m^2) where the aerial interpretation represents wetted channel but not the ground-based raster file; and (iii) the area (m^2) where the ground-based file represents wetted channel but not the aerial interpretation. Results are summarized in Table III and a map of the (ii) and (iii) intersections is presented in Figure 5e. The values in Table III yield an F statistic of $F = \frac{19289}{22450} = 0.859$, 85.9%. Thus, considering the ground-based outline as the true, there is a 14% of difference in the wetted channel area calculated from the aerial interpretation although this error represents the combination of errors attributed to the image registration and to the water outline digitalization as discussed above.

Discussion

The comparison of the delineation of wetted areas by photographic interpretation versus field delineation (RTK GPS measurements) revealed differences, which reflect more than just errors due to the photo registration, mosaicing process, and the precision to connect RTK GPS field measurements. These include operator interpretation errors and generalizations when digitizing the water outline from the registered photos, as well as operator and interpretation errors in the field by the RTK GPS surveyors. Although the point accuracy of the individual RTK GPS measurements (c. 5 cm) is likely to exceed the accuracy of the registered images, the operator's interpretation might obscure this gain. There may also have been a slight change in flows throughout the course of the day contributing to subtle discrepancies in wetted area at the time of the RTK GPS survey and the time of the blimp survey. Unfortunately, for the experiment reported here there is no reliable way to decompose these various components of uncertainty from the overall differences produced between the two methods. However, it is reassuring to note that there is only a 14% overall areal difference in wetted area and that the overall pattern is entirely coherent (Figure 5e).

Upon closer inspection of Figure 5e, the black and grey areas (blue and red in colour version) generally mirror each other on opposite sides of the bank all along the wetted channel outlines. That is, for example, if an area on the right bank was wetted on the aerial photo, but not the ground-based analysis; then it is generally matched by an area on the left bank that was wetted in the ground-based analysis, but not the aerial photo. This is suggestive of an aerial photo that is scaled correctly but has been shifted slightly during registration. Indeed, this trend can be seen to switch at photo boundaries in the mosaic by comparing Figures 5c and e. Table III shows that the blue and red areas are within 5% of equaling each other, suggesting that the majority of the discrepancies are in balance and likely due to slight shift errors in the registration for individual photos in the mosaic.

Figure 6a shows a mosaic of 45 images across a respectably sized reach. For contextual purposes and mapping of in-channel and floodplain habitat patches and mosaics across the entire reach in the Feshie in 2007, this sort of rapidly acquired imagery is perfectly suitable. A closer look at the imagery in Figure 6b and c reveals some minor misalignment of linear fluvial features between images due to registration errors in the mosaiced photos (largely

Table III. Wetted channel outlines matrix: aerial interpretation and ground-based outlines

Image interpretation	Ground-based	Wet Area (m^2)	Dry Area (m^2)
Wet Area (m^2)		19289[1]	1618[2]
Dry Area (m^2)		1543[3]	—

[1] Area (m^2) where both files represent the wetted channel.
[2] Area (m^2) where the aerial interpretation represents the wetted channel but not the ground-based raster file.
[3] Area (m^2) where the ground-based file represents the wetted channel but not the aerial interpretation.

down to inadequate GCP density). Such errors are generally on the order of 0.1–1 m. These can be improved somewhat by selecting images with more GCPs present, or by adding tie points in the registration based on easily distinguishable features that were surveyed in the field.

Part of the appeal of simple, low-cost aerial platforms like blimps for acquiring aerial imagery is the potential to undertake repeat surveys at higher frequency. If the imagery is simply used for contextual purposes to aid in interpretation of other datasets collected concurrently (e.g. topographic surveys) the magnitude of positional inaccuracies in the images are unlikely to be too problematic. However, the variable exposure within the mosaiced photos probably makes them unsuitable for most automated image classification applications. Even if manual digitization of vector line work is used to classify such imagery, the discrepancy in wetted channel outlines highlighted in Figure 5e highlights that caution should be exercised in interpreting differences between images such as lateral plan form changes (e.g. bank erosion and point bar development). However, if this information is combined with other evidence collected concurrently (e.g. field measurements or high resolution topography), then sensible interpretations can probably be inferred.

CONCLUSIONS

This study examined two field experiments to evaluate the field and laboratory methods necessary to generate georeferenced photo-mosaics from aerial photographs taken using a digital camera mounted on a helium blimp and provide an assessment of their reliability. Photo-mosaics can be used easily to facilitate routine systems monitoring, for example mapping dynamic channel morphology or changes in habitat suitability. Analysis of the local image texture can also be used to automate mapping of fluvial substrate size, capable of resolving grain-, cluster- and patch-scale morphologies. The main conclusions about the acquisition of aerial photography from such versatile, low-cost platforms as used in these experiments are as follows:

1. At high GCP densities, higher order polynomial transformation models provide the highest quality registrations, producing imagery suitable for applications such as automated grain-size detection. However, at more modest GCP densities (i.e. 19–28 GCPs Ha^{-1}), simple aerotriangulation and 2nd order polynomial transformation models perform better, resulting in registration errors at standards equal to or better than obtained with conventional aerial photography (e.g. 0.5–1 m).
2. The quality of image registration is highly dependent on the configuration of GCP targets. A configuration that spreads the coverage throughout the photograph undergoing registration, while marinating sufficient density is best. GCP RMSEs are minimized when tie-points are spatially concentrated; whereas RMSEs increase substantially in areas on the periphery of the GCP network.
3. Reach scale blimp surveys can be registered using relatively low GCPs densities (e.g. 9 GCPs Ha^{-1}) by means of applying an aerotriangulation transformation. Image processing errors vary throughout the image, and depend on the proximity to GCPs although errors are rarely found to exceed 1 m (a value consistent with the F statistic obtained for the wetted channel outlines analysis presented in the 'Results' Section).
4. Mosaiced aerial photographs of an entire reach at 5–10 cm resolution, and of adequate quality to support mapping of river corridors, can be acquired by low-cost blimp surveys. In the example reported here, the survey required less than a single day field work and laboratory processing, and presents a cost-effective alternative to traditionally commissioned flights. Such low-cost and time-efficient methods might make short time scale repeat surveys a more viable monitoring methodology in geomorphology, ecology and enginery studies.

REFERENCES

Booth DT, Cox SE, Simonds G. 2007. Riparian monitoring using 2-cm GSD aerial photography. *Ecological Indicators* **7**(3): 636–648.

Brasington J, Langham J, Rumsby B. 2003. Methodological sensitivity of morphometric estimates of coarse fluvial sediment transport. *Geomorphology* **53**(3–4): 299–316.

Brasington J, Smart RMA. 2003. Close range digital photogrammetric analysis of experimental drainage basin evolution. *Earth Surface Processes and Landforms* **28**: 231–247.

Carbonneau PE. 2005. The threshold effect of image resolution on image-based automated grain size mapping in fluvial environments. *Earth Surface Processes and Landforms* **30**(13): 1687–1693.

Carbonneau PE, Bergeron N, Lane SN. 2005. Automated grain size measurements from airborne remote sensing for long profile measurements of fluvial grain sizes. *Water Resources Research* **41**(11): W11426.

Chen X, Vierling L. 2006. Spectral mixture analyses of hyperspectral data acquired using a tethered balloon. *Remote Sensing of Environment* **103**(3): 338–350.

Church M, Hassan MA, Wolcott JF. 1998. Stabilizing self-organized structures in gravel-bed stream channels: field and experimental observations. *Water Resources Research* **34**(11): 3169–3179.

Fonstad MA, Marcus WA. 2005. Remote sensing of stream depths with hydraulically assisted bathymetry (HAB) models. *Geomorphology* **72**(1–4): 320–339.

Hodge RA, Brasington J, Richards KS. 2008. The application of terrestrial laser scanning to the measurement of grain-scale fluvial morphology. *Earth Surface Processes and Landforms* (accepted).

Hughes ML, McDowell PF, Marcus WA. 2006. Accuracy assessment of georectified aerial photographs: implications for measuring lateral channel movement in a GIS. *Geomorphology* **74**: 1–16.

Lejot J, Delacourt C, Piegay H, Fournier T, Tremelo M. 2007. Very high spatial resolution imagery for channel bathymetry and topography from an unmanned mapping controlled platform. *Earth Surface Processes and Landforms* **32**(11): 1705–1725.

Lillisand TM, Kiefer RW. 2000. *Remote Sensing and Image Interpretation (4th Edition)*. Wiley: New York; 736p.

Marcus WA, Fonstad MA. 2008. Optical remote mapping of rivers at sub-meter resolutions and watershed extents. *Earth Surface Processes and Landforms* **33**(1): 4–424.

McMillan HK, Brasington J. 2008. End-to-end flood risk assessment: a coupled model cascade with uncertainty estimation. *Water Resources Research* **44**: W03419. DOI:10.1029/2007WR005995

Mertes LAK. 2002. Remote sensing of riverine landscapes. *Freshwater Biology* **47**(4): 799–816.

Surian N, Cisotto A. 2007. Channel adjustments, bedload transport and sediment sources in a gravel-bed river, Brenta River, Italy. *Earth Surface Processes and Landforms* **32**: 1641–1656.

Verdu JM, Batalla RJ, Martinez-Cassasnovas JA. 2003. Estimating grain size distributions of a gravel riverbed at reach scale from detailed aerial photos, geostatistics and digital image processing (Isabena River, Spain), *Paper presented at the Braided Rivers Conference,* Br. Geomorphol. Res. Group, London, April.

Verdu JM, Batalla RJ, Martinez-Casasnovas JA. 2005. High-resolution grain-size characterisation of gravel bars using imagery analysis and geostatistics. *Geomorphology* **72**(1–4): 73–93.

Vetrella S, Tripodi C, Colagiovanni C, Alfano A. 1977. Tethered balloons as geostationary platforms for multispectral radiometry. *Acta Astronautica* **4**(5–6): 617–624.

Vierling LA, Fersdahl M, Chen XX, Li ZP, Zimmerman P. 2006. The short wave aerostat-mounted imager (swami): a novel platform for acquiring remotely sensed data from a tethered balloon. *Remote Sensing of Environment* **103**(3): 255–264.

Werritty A, Ferguson R. 1980. Pattern changes in a Scottish braided river over 1, 30 and 200 years. In: *Timescales in Geomorphology,* Cullingford RA, Davidson DA, Lewin J (eds.). Wiley: Chichester; 53–68.

Wolf PR. 1984. *Elements of Photogrammetry.* McGraw-Hill: New York; 624p.