THE PHOTOGRAMMETRIC RECORD

The Photogrammetric Record 32(160): 398–423 (December 2017)
DOI: 10.1111/phor.12213

MAPPING FIELD PHOTOGRAPHS TO TEXTURED SURFACE MESHES DIRECTLY ON MOBILE DEVICES

CHRISTIAN KEHL* (christian.kehl@uni.no)
SIMON J. BUCKLEY (simon.buckley@uni.no)
Uni Research AS, and University of Bergen, Bergen, Norway

SOPHIE VISEUR (viseur@cerege.fr)
Aix Marseille Université, Marseille, France

ROBERT L. GAWTHORPE (rob.gawthorpe@uib.no)
University of Bergen, Bergen, Norway

JAMES R. MULLINS (james.mullins@abdn.ac.uk)
JOHN A. HOWELL (john.howell@abdn.ac.uk)
University of Aberdeen, Aberdeen, UK

**Corresponding author*

Abstract

The mapping of photographs to surface geometry is an important procedure for many applications within the geosciences. This paper proposes an interactive framework for feature-based image-to-geometry mapping that works directly on mobile devices, under challenging imaging conditions and with limited available hardware performance. The framework makes use of openly available digital elevation models (DEMs) together with mobile position-and-orientation sensor data. It integrates calculation heuristics for result evaluation and feedback, synthesising available knowledge in current registration literature. The approach is assessed on two image datasets captured on separate occasions. Their interpretations are mapped to one textured lidar surface model and the projection accuracy is qualitatively assessed. The experiments show a significant accuracy improvement in photograph registration results, as well as the faithful mapping of image interpretations on the underlying surface geometry. This semi-automatic, user-guided, interactive approach is superior to comparable fully automatic registration methods.

KEYWORDS: feature-based registration, geological interpretations, image-to-geometry, interactive framework, mobile-device application, outdoor environments

Introduction

Automatically registering arbitrary, outdoor-landscape photographs to existing coloured geometry is beneficial to a large array of applications in different disciplines, in particular cultural heritage (Corsini et al., 2013; Pintus and Gobbetti, 2015), painting preservation (Remondino et al., 2011) and outcrop studies via multispectral imaging (Sima and Buckley, 2013), and geosciences visualisation and interpretation (Sima et al., 2013; Mullins et al., 2016). In order to extract a faithful and reliable registration for a given application demands addressing challenges such as the provision of initial positioning, the correlation of 2D-image and 3D-surface information, the treatment of perspective and radiometric differences, and the accessible, user-friendly parameterisation of such automatic techniques. The need for visual registration has increased recently due to its various applications, notably the re-texturing of cultural heritage models with uniformly lit photographic information (as shown by Dellepiane et al., 2012), reliable 3D outdoor navigation, the tracking of waterlines in river-stream flooding events (as presented by Kröhnert, 2016), the addition of semantic information to urban models and the annotation and interpretation of geological outcrop studies during fieldwork using mobile devices such as smartphones and tablets (Viseur et al., 2014; Kehl et al., 2015).

Desktop tools for automatic image-to-geometry registration have been available for several years. Recent developments allow this registration task to be performed on mobile devices, as presented by Gauglitz et al. (2014) and Kehl et al. (2015), though with significant accuracy limitations. The techniques provided by desktop tools are able to robustly and accurately register photographs to surface geometry, such as *mutual information* (MI) approaches. The implementation of these methods on mobile devices is technically problematic. In addition, the most successful desktop approaches commonly demand a minimum degree of user input for initial 3D positioning and orientation, as present in software applications such as MeshLab (Callieri et al., 2003). Mobile-device approaches aim to replace 3D user input with the acquisition of geolocation sensors. This results in georeferenced, metrically scaled, fully automatic approximation of captured images' external orientations. The acquired sensor-based external orientation via consumer-grade geopositioning-system (GNSS) measurements and a magnetometer is subject to a high degree of noise and inaccuracy, as discussed by Blum et al. (2012) and Kehl et al. (2015).

This paper introduces a novel framework for registering outdoor photographs to textured surface geometry in the field using mobile devices. Experiments and published results (Dellepiane et al., 2012; Corsini et al., 2013; Kehl et al., 2015) show the limitations of fully automatic approaches without user intervention. Hence, this paper proposes an interactive, visual approach within which the user can steer the application via digital elevation model (DEM) repositioning, image processing and parameter adaptation. A visual feedback loop provides multiple, visual estimates of the registration accuracy upon which the user can simply refine each image registration individually.

The envisaged workflow could potentially follow the following steps:

(1) Given that the domain expert acquires a photograph of a target object, an initial coarse external orientation is derived for the photograph using mobile sensor data via a magnetometer, global navigation satellite system (GNSS, for latitude and longitude, see Kehl et al. (2016)) and a coarse DEM (for the related altitude).
(2) The domain expert runs an initial, automatic feature-based image-to-geometry procedure (for example, using pure scale-invariant feature transform (SIFT) by Lowe (2004) and random sample consensus (RANSAC) EP*n*P (Torr and Zisserman, 2000; Lepetit et al., 2009)).

(3) The user is presented with a preview of the resulting registration that allows comparing the target photograph and the rendered (*synthetic*) image.
(4) Under realistic lighting conditions, the synthetic image will not match the photograph so that the user needs to adapt the image processing, the initial sensor-derived external orientation or the parameterisation. This refinement is supported by visual feedback cues on multiple levels.
(5) After a limited number of iterative refinement steps, the synthetic image resembles the photograph so that the domain expert can accept the computed pose as being "correct". The assessment is further supported by global quality criteria and their visual representation.
(6) The interpretations created on the acquired image can be accurately projected onto the textured surface geometry in 3D.

The primary intended application of such a workflow is the image-based surface interpretation of geological outcrops. The framework introduces mobile-device implementations of recent techniques for illumination-invariant registration methods suitable for outdoor imagery. Furthermore, the resulting implementation facilitates simple interfaces for mobile sensor data correction and a clear communication of the expected result accuracy. The user is guided towards optimal registration results via step-by-step evaluation and textural suggestions for process intervention. The approach is assessed using a geological interpretation scenario of Mam Tor in Derbyshire, UK.

RELATED WORK

The available literature related to this paper's background has increased substantially in recent years. A growing number of research groups within the geosciences envisage integrating mobile devices and 3D lidar surface data into their workflows (McCaffrey et al., 2005). Applications can be found across the geosciences, whereas the context of this research is in the supported interpretation of outcrops for geomodelling purposes (Enge et al., 2007; Howell et al., 2014) and virtual fieldtrips for geological training purposes (McCaffrey et al., 2010). This section distinguishes between: (1) recent advances in image-to-geometry registration procedures; and (2) state-of-the-art geosciences applications using mobile-device technology, with a specific focus on those using 3D data.

Advances in Image-to-geometry Registration

A large array of computer graphics, computer vision and photogrammetric desktop software is available registering photographs on an available textured surface model in 3D. Many available desktop tools use techniques such as *spatial resection* (see Moffitt and Mikhail, 1980; implemented in, for example, Riegl RiSCAN), *mutual information* (MI, see Maes et al., 1997, implemented in MeshLab) or *structure-from-motion* (SfM; see Snavely et al., 2006, implemented in Photo Tourism (GRAIL, 2017)). MI in MeshLab (introduced by Corsini et al., 2009 and subsequently refined by Sottile et al., 2010) is an intuitive tool for 2D–3D registration. The unique ability to also register images to untextured geometry is particularly advantageous. The SfM registration approach has seen a widespread adoption in recent desktop applications (Corsini et al., 2013; Pintus and Gobbetti, 2015). Both approaches have shown their potential and capabilities in several desktop application scenarios, but their adoption on mobile devices is currently prohibitive due to technical obstacles. As promising as MI algorithms appear, they are rarely ported to mobile-device

platforms due to their reliance on specific mathematical system libraries, such as LAPACK, BLAS and NEWUOA (Powell, 2006). SfM approaches yield high registration accuracies, but the involved computations currently prohibit mobile-device implementation, where hardware and desired calculation times are limited. Furthermore, SfM demands a large, overlapping-image collection to be available. This key demand contrasts with actual cases that rely on single photographs or image sets containing only a few pictures, mapping individual image interpretations on the geometry.

Feature-based registration is the most prevalent technique for image-to-geometry registration on mobile devices, as used by multiple studies in recent years (Gauglitz et al., 2011b; Kehl et al., 2015; Sweeny et al., 2015; Kröhnert, 2016). A recent study presents technical advances for treating illumination differences in outdoor environments – a long-standing challenge for feature-based registration (Kehl et al., 2017). Other approaches have shown that real-time registration is technically feasible (Gauglitz et al., 2014), which will open up new application possibilities in the future.

State-of-the-art of Mobile Applications within the Geosciences

Mobile devices are increasingly popular tools within the geosciences to study natural phenomena outdoors. Base-mapping software (such as Google Maps or OpenStreetMap) is ubiquitously available and readily accessible in urban and densely populated areas. Although map sections can be buffered on a mobile device to be accessible without web connectivity, the usability of base mapping in sparsely or unpopulated areas of the globe is limited. This is because the general lack of base-mapping detail and the lack of web connectivity to load unbuffered areas or further detail levels. Furthermore, data representation is often 2·5D (where each X, Y point has only one Z value) which has limited applicability for tasks outside geographic orientation and aerial mapping.

Mobile applications that facilitate particular tasks within field-based geosciences and outcrop geology have recently appeared. The Midland Valley (2017) "Fieldmove Clino" is a commercial application that allows the logging of positions and strike-dip orientations, the storage of georeferenced globally positioned photographs and notes, and digital field mapping and editing on orthophotographs. Systems for georeferenced interpretation and measurement registration have been previously presented (for example, Dey and Ghosh, 2008). A recent review of mobile-device tools showed applications that allow keeping track of geological measurements and notes (Ferster and Coops, 2013) in the form of a "digital geological fieldbook". Georeferencing in related tools is limited to global latitude–longitude coordinates. Additionally, commercial drawing applications (such as Autodesk SketchBook or NVIDIA Dabbler) for mobile devices are intuitive and accessible for most novice users to draw, sketch and annotate acquired field photographs. On the other hand, these applications lack 3D registration, a georeferencing component or any support to enrich the annotations with supplemental information at a later stage.

Most early approaches lacked the 3D component necessary for "vertical geology" to be captured, analysed, interpreted and stored. Recent sedimentary logging applications, such as SedMob (Wolniewicz, 2014) and Strataledge from Endeeper (2016), now provide such functionality. The subsequent registration of 1D vertical data (for example, facies logs) to outcrop surface models in 3D is currently a manual process; however, this could be done automatically using the framework presented in this paper.

Considering applications in the geosciences and the geoinformation sciences that already make use of registered information in 3D space using mobile devices, a small number of academic prototypes are available in the domain. 3D data on handheld devices is

used for scene understanding (Wang et al., 2012), urban environments and tourism (Schilling et al., 2005), and geo-navigation on a mobile device. A mobile application for local flooding mitigation (Kröhnert, 2016) relies on accurate image-to-image and image-to-geometry registration. A milestone application in geology, called "Outcrop", allows the interpretation of digital outcrop surface models in 3D (Viseur et al., 2014).

In conclusion, one of the largest issues for better 3D data utilisation is the gap between desktop and mobile software tools and their facilitated features. Despite this gap, mobile devices offer opportunities for 3D registration and referencing by querying built-in position – orientation sensors. This ideally replaces the need to involve 3D initial estimates via user interaction, if the sensor accuracy can be improved. Current position sensors (GNSS) in mobile devices are not able to provide sufficiently reliable and accurate position measurements (Kehl et al. 2016). Although the sensor inaccuracies are well known, the available tools do not facilitate a user-based correction of the sensor results, nor an accessible way of user intervention in the mobile registration. The framework presented in this paper specifically addresses these challenges and issues.

METHOD

The adopted approach for the image-to-geometry registration in this paper aligns with the workflow previously introduced in Kehl et al. (2016). In that approach, the system is provided with the textured surface geometry, its georeferenced centre point and the camera configuration of a specific mobile-device camera that is used for the image acquisition. Then, field photographs are taken with the mobile device's camera and the location and orientation sensors are logged at the time of image capture. In contrast to the initial workflows, an optional DEM-based altitude estimation is adopted in this new workflow, where the altitude is interpolated on a coarse grid which is provided a priori. Subsequently, a virtual camera is positioned within the scene with respect to the surface geometry by using the logged location and orientation, which allows capturing a corresponding rendered image (previously referred to as a *synthetic image*; see Kehl et al. (2015, 2016) for the terminology). A feature-based registration approach computes the point correlations between both images, intersects the 2D features with the surface geometry in 3D and calculates a detailed exterior orientation using a *perspective-n-point* (P*n*P; Lepetit et al., 2009) method. This detailed exterior orientation allows a projection of image-based interpretations onto the surface geometry. In the initial implementation, the registration quality of the computation was largely unknown and not communicated to the user. User feedback and intervention were kept to a minimum, even though some intervention is needed to improve the registration. In particular, changing environmental conditions, outdoor illumination and poor positioning estimates could lead to registration failure in many usage scenarios.

In this paper the authors explain how to implement effective user feedback and intervention into the previously reported method. The workflow is thus adapted within the following stages (Fig. 1):

(1) In the sensor logging stage, the adapted workflow allows the user to interactively correct the sensor position both laterally and vertically by offsets using the DEM, and to interactively render the scene from user-defined camera locations.
(2) After capturing the photograph and rendering the reference image, both images are optionally subject to image-processing procedures to counter problematic radiometric influences (Kehl et al., 2017) in the feature matching.

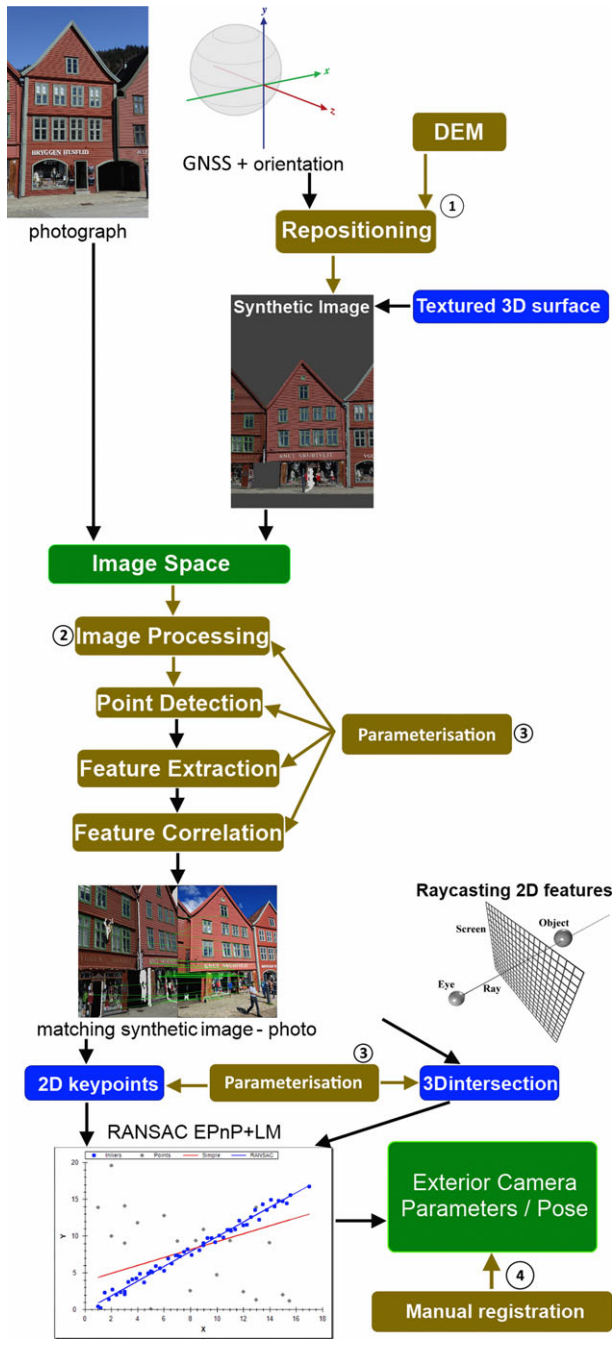

FIG. 1. Adapted workflow from Kehl et al. (2015), showing possibilities of user intervention on the registration results in various stages, as indicated by the brown boxes. Circled numbers refer to the four stages in the text.

(3) Apart from algorithm-specific parameterisation, there are common parameters for the feature matching that have a greater or lesser significant influence on the registration (Mikolajczyk et al., 2005; Gauglitz et al., 2011b; Sima and Buckley, 2013; Kehl et al., 2016), which can be adapted and iteratively refined.
(4) In case the semi-automatic registration method fails, it is possible to manually and interactively generate conjugate feature points on the mobile device using the given textured surface (also introduced in Kehl et al., 2016) and captured photographs.

Feedback is given to the user in multiple ways: the relocation interactively provides the user with an updated rendered image. For the parameterisation of stages (2) and (3), quality metric statistics are tracked over subsequent registration trails, hence indicating to the user a metric change and a metric's absolute qualification (background colours indicating "superior", "normal" and "inferior" metric results). Based on previously researched quality metrics and registration heuristics (Kehl et al., 2017), a textural hint for parameterisation adaptation is given to guide the user. An overall quality evaluation is given in a five-stage colour range of the screen background to allow simple quality judgements by the user for the termination of the process.

The following subsections present further details about the DEM repositioning (1), the image processing (2), the reference image capture (4), and the statistical evaluation, communication and parameter adaptation (3).

DEM Altitude Adjustment

The estimated position of a captured photograph within the previously presented approaches is fixed, as it is determined automatically by sensors upon photograph acquisition. The statically inferred altitude (for example, via GNSS) is significantly influenced by signal noise and reception, as discussed in previous field studies (Kehl et al., 2016). The altitude can potentially be tracked over longer periods, but the prospective signal improvements are too marginal, in relation to power consumption costs, to make consumer GNSS or barometric sensors a viable positioning source. High-end multi-frequency and differential GNSS (DGNSS) allows highly accurate positioning, but the licensing and usage costs make this solution infeasible in standard cases (for example, with ad-hoc fieldwork or educational group excursions). In this paper altitude estimation via a supplementary DEM is presented as a viable alternative to GNSS altitude measurements. In a first stage, the process allows just static positioning with a local DEM provided for interpolation. A statistical evaluation in Fig. 2 shows that the average altitude error (in terms of reference-to-result absolute distance value) decreased by the application of a $25\,\mathrm{m} \times 25\,\mathrm{m}$ DEM, but the standard deviation of the altitude error increased. The incorporated data are the 13 correctly registered images, out of a total of 17, within the Mam Tor September image set, presented in the "Dataset" section. This increase is potentially due to the coarse lateral DEM resolution, which leads to an inaccurate DEM gradient in steep terrain sections. Stochastically sampled local lateral adaptations potentially present further automatically calculable position improvements. On the other hand, such an approach means rendering and computing the pipeline in Fig. 1 for multiple alternatives, which is computationally infeasible for current mobile-device platforms.

In response to this improvement, the user is presented with an option to locally adapt the initial altitude. This paper proposes using the DEM provided, instead of available base mapping (such as Google Maps or OpenStreetMap). Modern base-mapping systems allow

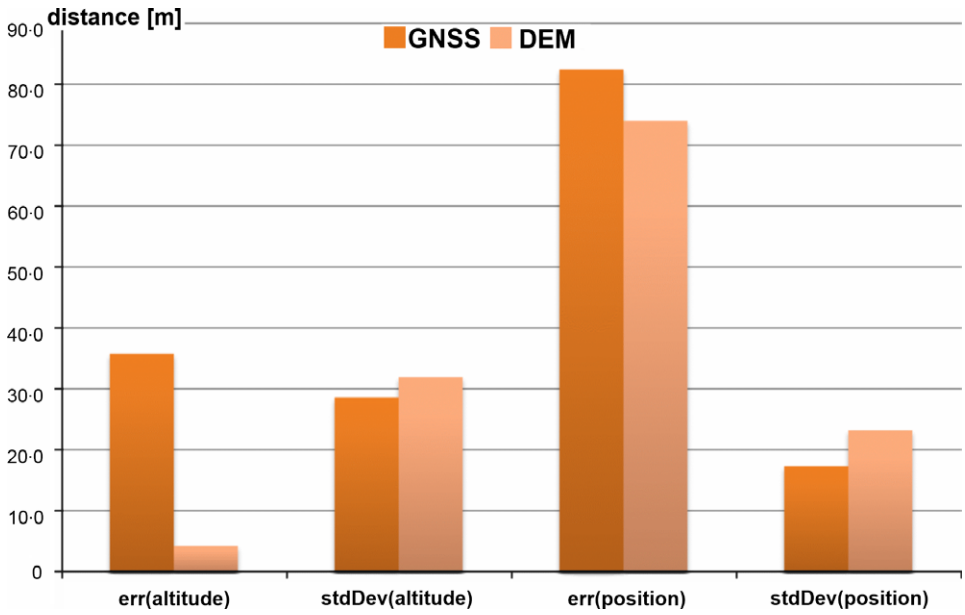

FIG. 2. Accuracy assessment of positional improvements (in metres, right-hand bars) – separately also assessed for the altitude contribution (left-hand bars) – using DEMs for altitude estimation. The DEM altitude is more accurate than GNSS altitude, but it is also affected by a higher standard deviation.

visual relocation and location querying within simple, intuitive and convenient user interfaces. A stable web connection or largely buffered data are mandatory, which is not feasible for outdoor applications. The DEM used for the presented test has a file size of less than 50 megabytes, which is rapidly loaded, buffered and queried for new locations in the field. A simple user interface for zooming, undoing, rescaling and confirmation make the altitude modification accessible to the user, as shown in Fig. 3(b). Linked render-view updates are based on newly picked DEM positions so that the user has direct visual feedback on the updated exterior orientation, as shown in Fig. 3(c).

Image Processing and Adaptations to Radiometric Influences

As discussed in the section "Related Work", the treatment of radiometric influences (such as illumination changes, moisture and environmental effects) between a photograph and a reference rendered image is challenging and, based on the state-of-the-art in computer-vision research, cannot currently be resolved automatically without full, 3D lighting reconstruction. Radiometric influences primarily affect the luminance channel of images. In feature-based image-to-geometry registration, point-feature algorithms operate on the luminance channel for establishing image-to-image correspondence, leading to high sensitivity to radiometric effects.

Recent literature describes the application of different point-feature algorithms to establish image correspondences. Key results of the related study in Kehl et al. (2017) present two suggested algorithms ((i) maximally stable colour regions (MSCR; Forssén, 2007); and (ii) features from accelerated segment test (FAST; Rosten et al., 2010)) that

(a)

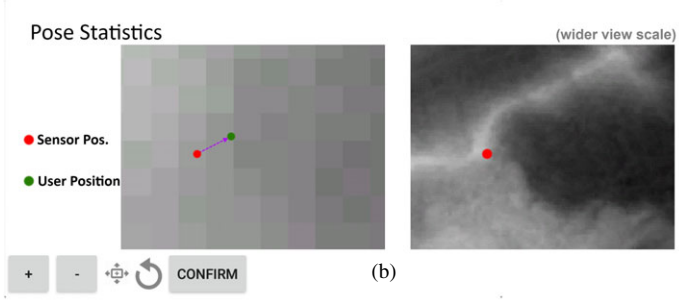

(b)

(c)

FIG. 3. Interactive DEM repositioning. (a) An incorrect, initial exterior orientation is created via static DEM look-up. (b) The user relocates the viewpoint of the camera with a simple, map-like user interface. (c) The adapted, improved external orientation is subsequently communicated to the user.

regularly outperform the standard registration based on the SIFT (Lowe, 2004). Additionally, both algorithms deliver complementary results (that is, they seem to succeed in opposite cases of radiometric influence). Therefore, the presented framework allows the user to select different algorithmic combinations, allowing adaptations to unsatisfying registration trials. The side-by-side image view and the correspondences image (see Fig. 4, which indicates point-to-point line connections) provide the user with visual information

about the registration process. The registration framework makes use of OpenCV4Android (a mobile version of OpenCV; Bradski and Kaehler, 2008) and osgAndroid (a mobile version of OpenSceneGraph for 3D rendering).

Another feature is to allow simple image processing within the registration pipeline because linear (brightness-contrast) and non-linear filters (Wallis, 1974) can directly influence the luminance channel. This, in return, directly influences the registration process, as previously studied by Jazayeri and Fraser (2010). The user interface provides a choice of

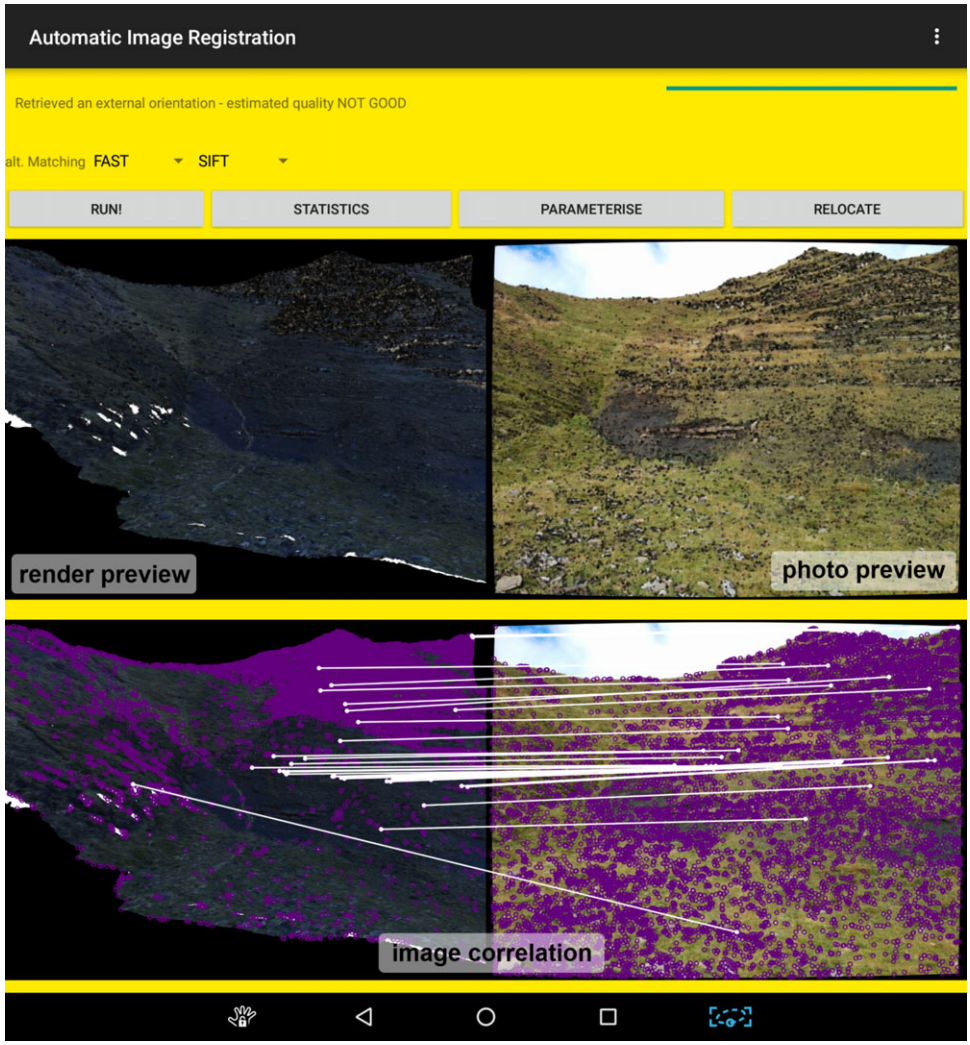

FIG. 4. The mobile-device user interface, as presented to the geoscientist. A choice of feature detectors and descriptors is given. Possibilities for statistical evaluation, registration parameterisation and repositioning are available on command. Before the registration, previews of the photograph (centre right) and rendered image (centre left) are presented, whereas the point correlation (bottom) is shown after the process.

algorithms that can be applied, together with their corresponding parameters. As the impact of the processing can hardly be determined a priori, the application or omission of image processing is decided by the user.

Statistical Evaluation and Communication

At the highest level of quality communication, a five-grade classification is used to communicate the overall quality. For a given registration, each metric is evaluated as being optimal (+1), normal (0) or insufficient (–1). All metric values P, together with their quality weight w, are composed into a quality measure Q as follows:

$$Q = \sum_{i=1}^{7} P_i \cdot w_i. \tag{1}$$

Fig. 5 shows an example qualification, where the legend shows the ranges and quality names for the five-scale classification system. This visually clear feedback is designed to give a novice, non-technical user quick and clear information about the success of the operations.

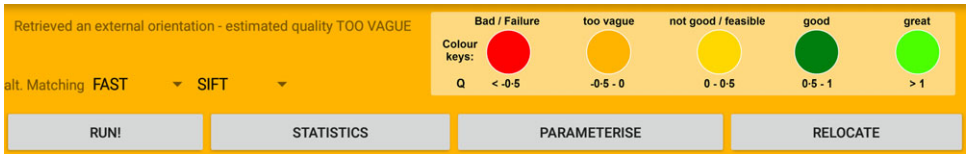

FIG. 5. Example for communicating overall quality via text (top-left) and background quality (top-right). The background quality is classified into five colour-coded "bins" with their quality metric limits (Q).

The quality assessment and the visual hints about the registration are based on a statistical evaluation of the feature correspondences and their 2D – 3D deviations. As discussed in Kehl et al. (2017), the pure measurement of point reprojection errors is an insufficient criterion to derive the registration quality. Thus, other metrics have been used, such as the rotation-quaternion and translation-vector differences (Kehl et al., 2016), or more expanded observations on homography (Quan, 2010) and least-squares constraints (maximum likelihood estimation sample consensus (MLESAC); Torr and Zisserman, 2000; Lowe, 2004) within the registration process, which are adopted in this paper. The following quality metrics, and their respective quality weight, are used in this paper's experiments:

(1) *#points* ("photo" and "render image"): the average number of points detected within each image; weight 0·1 per image.
(2) *#feature inliers*: the number of actual correspondence matches that adhere to epipolar geometry constraints, such as epipolar line distance and homography confidence; weight 0·25.
(3) *#feature correspondences in 3D*: a subset of feature inliers that have a valid 3D intersection with the reference surface geometry; weight 0·25.

(4) $r_{in(feat)}$: the ratio of feature inliers to the detected feature correspondences; weight 0·3.
(5) *# 3D feature inliers*: the number of 2D – 3D point pairs that are within the reprojection error limit after a RANSAC P*n*P pose estimation; weight 0·3.
(6) $r_{in(opt)}$: the ratio of 3D feature inliers to the feature correspondences in 3D; weight 0·3.
(7) *pixel reprojection error* [$\Delta(px_{feat}, px_{proj})$]: the reprojection error of the camera-projected 3D feature inliers to their 2D reference location in the photograph; weight 0·3.

These metrics are presented to the user in an overview, as well as the specific point-to-point mapping a posteriori, using an OpenCV match image (see Fig. 4). As each metric has a lower limit of its operability or feasibility, and because there are empirical values per metric indicating a lower optimum, the registration results can be simply communicated to the user in an accessible manner, as shown in Fig. 6. An experienced user can conclude parameterisation changes based on this feedback. On a higher level, a previous paper (Kehl et al., 2017) suggests a heuristic upon which to judge the registration quality and statistical metrics. This heuristic is applied in the presented framework to automatically deduce parameterisation changes. This logic is presented to the user in the form of an explanatory text, as shown at the bottom of Fig. 6. Furthermore, the heuristics even allow an automatic adaptation of the registration parameters for successive trials. The automatic evaluation heuristics, on which a textural hint is based, function according to the following rules (in descending priority order):

(1) *# salient points in photo or synth. image < 500*: too few points; adapt point detection metrics (lower thresholds, larger margins).
(2) *# correct feature correspondences [inliers 2D] < 7*: too few correspondences; erroneous point correlation; common error sources are drastic radiometric differences; adapt matching by image-processing filters.
(3) *# correct 3D features after pose estim. [inliers 3D] < 7*: RANSAC EP*n*P constraints very strict (too high a demand on reprojection accuracy); possibly arbitrarily distributed points (no match between photograph and synthetic image possible); adapt reprojection accuracy constraint.
(4) *feature correspondence ratio [inlier ratio (2D)] < 0·1*: again too few correspondences and erroneous correlations, possibly from radiometric variance; adapt matching by image-processing filters.
(5) *inlier ratio (2D) > 0·1 AND inlier ratio (3D) < 0·1*: the 2D feature point correlation is accurate and correct, but the 2D features are badly localised. This leads to a large reprojection error of their 3D intersections after RANSAC EP*n*P. It is suggested to slightly relax the correlation constraints (epipolar constraint; Torr and Zisserman (2000)) to increase the possibility for more 3D points in the pose optimisation.
(6) *inlier ratio (2D) > 0·1 AND inlier ratio (3D) > 0·1 AND avg. point dist > 2·0*: 2D point correlation is good and there are enough reprojection candidates to estimate a pose. The pose can be considered "correct", but the accuracy can be improved (due to the high reprojection error). In contrast to the previous case, it is suggested to impose stricter correlation constraints to reduce erroneous points in the EP*n*P pose estimation.

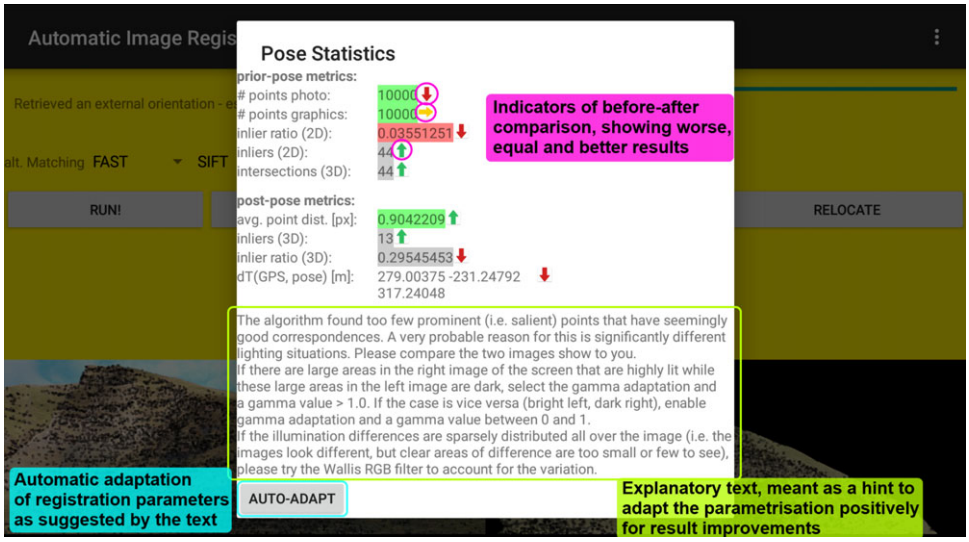

FIG. 6. In the statistical evaluations view, the overall quality of a single parameter is communicated via its background colour. The quality trend of successive trials is shown with coloured arrows next to each metric. Based on previously established limits and heuristics, an explanatory text (green highlight) and an auto-adaptation (cyan highlight) indicates the parameters adapted for new trials.

Creating Reference Images for Manual Registration

In some cases of extreme imaging conditions, significant illumination changes and insufficiently accurate input pose, it may be necessary to determine control points and 2D – 3D correspondences manually. The presented framework utilises the manual registration discussed in Kehl et al. (2015, 2016), where a major drawback was the mandatory, a priori provision of 3D-positioned keyframes with which to register 2D images. Because the 3D data are expected to be available on the mobile device in the field, it is possible to define easily recognisable surface control points on the spot by the user. Then, the virtual camera of the 3D-surface viewer can be freely positioned to interactively create the reference image, as shown in Fig. 7. These images are then input for manual registration using P*n*P pose estimation as reported in Kehl et al. (2016).

DATASETS

The studied datasets within this paper were acquired during two field campaigns to the Mam Tor turbidite outcrop in the Peak District, Derbyshire, UK. The study object is a cliff-shaped outcrop formation with mudstone – sandstone interbedding. The outcrop has been used to study multi-storey channel sandstone configurations in turbidites (Southern et al., 2014) for improved petroleum recovery (Pringle et al., 2006), and notably to study landslide processes (Waltham and Dixon, 2000) which presently shape and deform the outcrop.

In order to study 3D processes, terrestrial laser scans (TLS), combined with digital single-lens reflex (DSLR) imagery, were acquired in March 2015. The specific instruments used were a Riegl VZ-1000 lidar scanner supplement by a Nikon D800E camera with a Nikkor 85 mm lens. In accordance with the workflow described by Buckley et al. (2010)

FIG. 7. The integrated 3D model viewer ("Virtual Outcrop Renderer") of the framework allows defining 3D control points and capturing registration reference images (right, (a)). These reference images are used in a P*n*P-based manual registration process where the user selects 2D–2D point pairs (left, (b)).

and comparable to other digital outcrop surveying techniques (Bellian et al., 2005; Olariu et al., 2011; Richet et al., 2011; Hodgetts, 2013; Rarity et al., 2013; Howell et al., 2014), a textured surface mesh with a varying vertex sampling density has been created.

Simultaneously with the TLS and DSLR acquisition in March 2015, images were acquired using an LG Google Nexus 5 mobile camera (8 Mpixel) according to the workflow in Fig. 1. The photographic campaign resulted in a first reference dataset. During a second fieldtrip in September 2015, the photographic campaign was repeated, resulting in a second image set. The datasets differ in their radiometric properties because of varying illumination and higher moisture exposure of the rock within the September dataset. All datasets (namely, TLS-DSLR, March Nexus photographs and September Nexus photographs) are equivalent to previous studies of the subject by Kehl et al. (2015; 2016). In addition to this data, an openly available DEM with a laterally uniform resolution of 25 m was utilised for refining the initial exterior orientation values.

The March dataset was registered to surface geometry without significant issues. On that basis, a set of 2D geological interpretations of the outcrop were created, shown in Fig. 8(a). The practical aim of this paper is to illustrate that a similarly accurate 3D interpretation mapping can be achieved using the September dataset (Fig. 8(b)), in comparison with the reference March dataset. Previous studies have shown that the automatic registration of such a dataset does not succeed for a major subset of images because of the challenging imaging conditions. Therefore, by applying the presented registration framework, it is possible to improve the overall registration success rate to a level that allows valid mapping of image-based interpretations.

Mam Tor – March interpretations (a)

Mam Tor – September interpretations (b)

FIG. 8. Example image-based geological interpretations for the March 2015 (a) and September 2015 (b) datasets. The colour markers represent sandstone (yellow) and mudstone (grey) interbedding interpretations as well as major fault lines (red).

The interactive, visual approach is realised for mobile devices with the Google Android operating system. The particular device used for testing and experiments is an NVIDIA Shield tablet. This device was chosen for its high performance, dedicated graphics processing unit (GPU), beneficial computational properties and high-resolution screen (1920×1200 pixels on an 8-inch (20-cm) screen). The dedicated graphics core has an extended OpenGL instruction set that equates to common GPUs in desktop computers (in contrast to most other available mobile devices), which is particularly advantageous for the texture compression of the presented textured surface models.

RESULTS

A first focus of the experiments was the quantitative accuracy assessment of the presented interactive framework. The datasets of the March fieldwork (see Kehl et al., 2016) and the September expedition (see Kehl et al., 2017) with the DEM have been re-assessed for the interpretation, with a focus on a selection of crucially challenging images. These re-established quality statistics are compared with the registration quality achieved with the September dataset using the interactive framework (see Table I). The focus on crucial image sets is the reason for the reduced total number of assessed images in columns 3 and 4 of Table I.

TABLE I. Statistical comparison of achievable image-to-geometry registration results on the Mam Tor case study for the equal-condition March dataset, the poorly conditioned September GNSS dataset (both studied previously), the September dataset with static DEM positioning and the interactive approach on the September dataset.

Dataset/method	*Mam Tor March (corrected GNSS, Kehl et al. (2016))*	*Mam Tor Sept. (GNSS, Kehl et al. (2017))*	*Mam Tor Sept. (DEM)*	*Mam Tor Sept. (interactive framework)*
Algorithm[a]	SIFT–SIFT	SIFT–SIFT	FAST–SIFT	Various; manual
Image processing	None	None	None	None, gamma, Wallis
Avg. points	15 000	15 000	4396·5	2800·6
Avg. matches	5·86	0	10·83	3·94
Avg. $r_{in(feat)}$	0·15961	–	0·77778	0·19441
Avg. $\Delta(px_{feat}, px_{proj})$	308·85	–	3·9244	17·316
Avg. $r_{in(opt)}$	0·51135	–	0·025397	0·42066
Correctly registered images	21/27	0/58	2/17	10/17 (auto); 3 (manual)

[a]The algorithm row shows the *detection* followed by the *description*.

The interactive framework yields a significant improvement in image registration quality. This result is validated in Fig. 9, where two images and their registrations are compared for each experimental series. For the March dataset, the *closest-to-comparison* image is selected, which is the accuracy reference that utilised manually altitude-corrected GNSS data. The September scenarios use basic GNSS data (latitude – longitude) and the DEM altitude ("Sept. DEM" in Fig. 9) or the interactive framework ("Sept. Interact." in Fig. 9). In conclusion, the intermixture of algorithms, image-processing techniques and the repositioning of the data yield the largest improvements. This is visible from column 4 resembling accurately the camera's external orientation of the photograph in column 1 of Fig. 9. The diagrams in Fig. 10 show the number of images registered with a

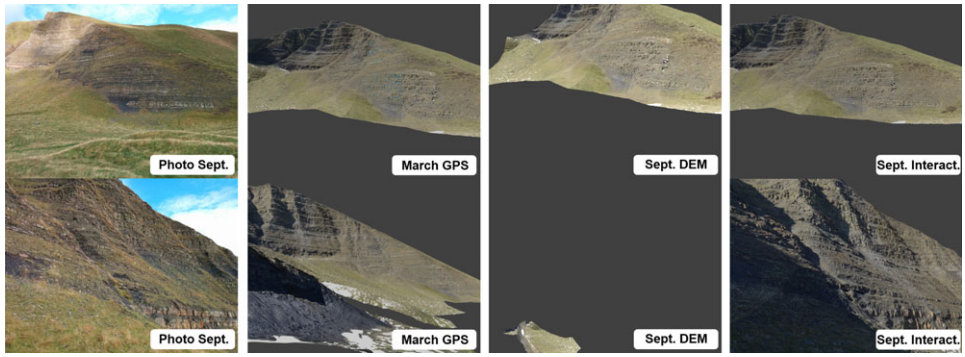

FIG. 9. Selective comparison of the visual registration quality of the presented approaches. The photograph from the March set is the reference to which the September data adhere. For the March comparison, the closest image of that selection is chosen for comparison, which is not identical in its position.

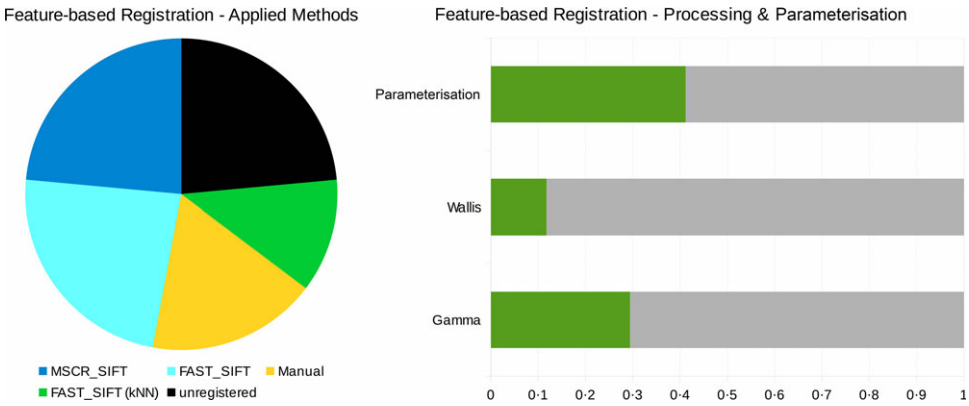

FIG. 10. Diagrams showing the statistical use of different feature matching methods (a) and process influences, such as method parameter changes and image processing (b), within the interactive registration framework when registering the Mam Tor dataset with DEM altitude initialisation from September 2015.

specific technique (a) or processed with a given filter (b), in relation to the total available images.

After the semi-automatic registration, the images without successful registration are registered manually. Fig. 11 provides an overview of the number of automatic (bright-green squares) and manual (dark-green squares) registrations within the static (a) and interactive (b) registration. A sensor pose (orange squares) represents failure with the available methods.

The interpretations have been projected on the surface geometry for a final, qualitative assessment. Fig. 12 (bottom) shows a panorama of the outcrop surface model with the registered image-based interpretations in the 3D scene along the cliff. The digital outcrop model (DOM) stays fixed in each case, while the image set and its registration is varied accordingly. The challenging image dataset in September (Fig. 12, middle and bottom) is compared to the reference interpretation of the March dataset (Fig. 12, top). Automatic registration without user intervention inadequately maps the interpretation to 3D, as seen in Fig. 12 (middle).

With respect to runtime performance, an update performance measurement of Kehl et al. (2015) is given in Table II. In this table, the mean runtimes and their standard deviation are split into the different stages of the registration process. The actual fine registration within the feature-based framework (referred to as *mutual correspondence* in the timings) occupies the major time spend in the method framework. Therefore, the process is further split into the 2D feature calculation and matching (that is, salient point detection, feature description and feature mapping) and the subsequent ray casting for retrieving 3D features. As can be seen when comparing the second and third columns of Table II with the fourth and fifth columns, the reduction of the model relates to only minor changes in the 2D feature calculation time (speed-up factor: 2 ×, due to the slightly reduced level of detail of the synthetic images). However, there is a drastic change in 3D-related operations, such as rendering the model into the synthetic image and the 3D ray casting (speed-up factor: 5× to 34×, due to reduced geometry to render or check for intersections).

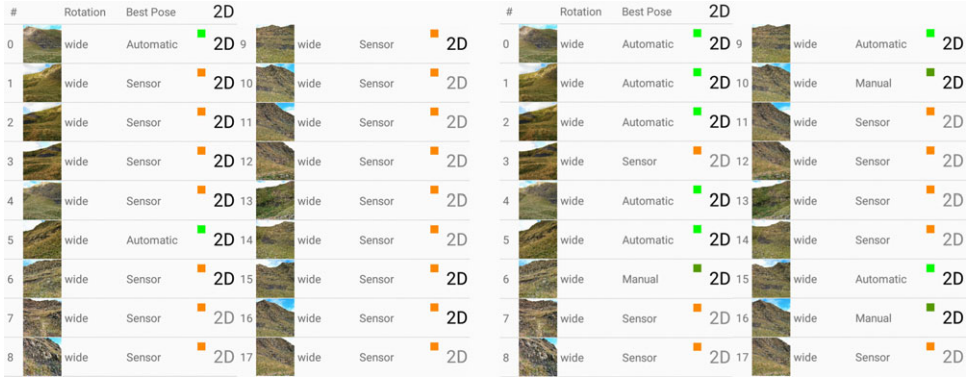

FIG. 11. Mam Tor September 2015 dataset overview, highlighting the number of successfully registered images (bright green for automatic registration; dark green for manual) in relation to the total number of images (orange (sensor) indicates registration failure). (a) The statically positioned DEM dataset. (b) The dataset processed interactively.

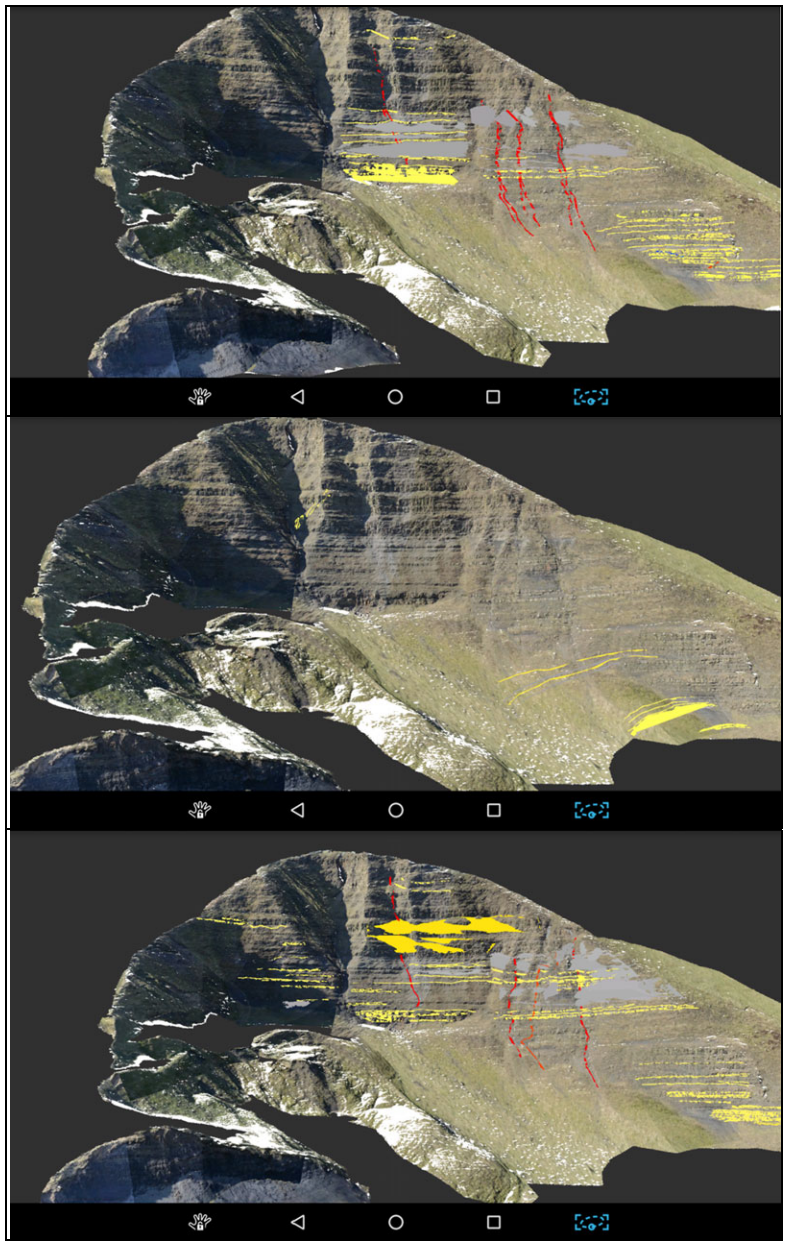

FIG. 12. Visual, qualitative comparison of the achievable interpretation mapping results from input, interpreted images (see Fig. 9 for interpretation of colours). The March dataset (top) was successfully registered as it was captured under the same imaging conditions as the 3D model, leading to high-accuracy interpretation mapping. If the image conditions change, the static registration (even with DEM support) is less stable, leading to errors in the interpretation mapping (September dataset, middle). This can be remedied using the presented, interactive framework (bottom), which allows creating interpretations with similar success as the reference mapping in March (top).

TABLE II. Processing runtime measurements for the registration process of the Mam Tor September dataset using the NVIDIA Shield K1 tablet (LTE version; 32 GB drive memory; 2 GB main memory; camera $f = 3{\cdot}92$ mm, $4{\cdot}6\,\text{mm} \times 3{\cdot}52\,\text{mm}$ CCD, 1632×1224 pixels). The timings (in seconds) are taken as the mean and standard deviation over 17 images (including partial registration failures) with varying image-processing strategies in place. The second and third columns relate to the current standard model configuration on the Tegra K1 tablets, containing around 1 200 000 triangles. The fourth and fifth columns relate to using a reduced DOM version applicable to tablets with lower graphical performance, for example, with Qualcomm Andreno GPUs, containing around 30 000 triangles and lower-resolution textures. Note that the computation time of the *mutual correspondence* method is further subdivided into its two main constituents (*detect and match*; *ray casting*) for an in-depth bottleneck analysis.

Metric	*Current standard model runtimes*		*Reduced DOM model runtimes*	
	Mean (s)	*Std. deviation (s)*	*Mean (s)*	*Std. deviation (s)*
Rendering preview	11·017	1·951	2·705	0·282
Load photograph	0·197	0·011	0·178	0·010
Compute coarse matrix	0·006	0·003	0·005	0·001
(Re-)rendering	10·757	2·021	2·537	0·173
Image processing	0·658	0·427	0·629	0·445
Mutual correspondence	172·098	98·154	37·720	18·510
Detect and match	*66·683*	*28·211*	*36·830*	*17·056*
Ray casting	*103·351*	*71·632*	*2·732*	*3·729*
Pose estimation	0·347	0·243	0·108	0·103

DISCUSSION

Technical Details

Introducing the DEM for the initial altitude determination improved the initial pose for the presented case study. Using the DEM for a manual refinement of the initial GNSS data has shown to be even more applicable because it facilitates a simple refinement without WiFi network connectivity. Despite the visual convenience and ease of use of the DEM, systematic errors were encountered during the experiments. Cell-boundary effects are large, due to the relatively low resolution of publicly available DEMs. This leads to cumbersome repositioning in steep and irregular terrain, particularly at positions close to vertical sections. This issue limits the potential for DEM-based repositioning in outcrop geology, where most fieldwork and the application of mobile devices is conducted close to vertical cliff sections. Additionally, altitude estimation accuracy is affected by general GNSS positioning errors, as the horizontal accuracy affects the altitude interpolation from the DEM. A significant issue to consider for practical use of DEM-based positioning in further field studies is its dependency on publicly available maps. The availability issue has been partially mitigated by recent Shuttle Radar Topography Mission (SRTM) observation and their distribution via the "Digital Earth Explorer".

Another technical aspect covered in this paper is the dependency of outdoor exterior orientation estimation on the underlying feature detection and correlation algorithms. This dependency, previously discussed by Mikolajczyk et al. (2005), Jazayeri and Fraser (2010), Gauglitz et al. (2011b) and Kehl et al. (2016), is confirmed by the presented results: the interactive registration framework returns significantly different results depending on the feature-detector and description-matching algorithm combination. Fig. 10 shows that the best

registration results have been achieved by mixing different algorithms and adapting their parameterisation specifically to each image. The applied image processing, such as Wallis filtering and gamma adaptation, also influences the registration result in the presence of illumination differences. In conclusion, Table III shows, in descending order, the key influencing factors that can be adapted by the user for achieving an optimal registration result. It has to be noted that the theoretically best registration on desktop computers could be achieved within the MI framework, as recently demonstrated by Guislain et al. (2017) for urban lidar data. The realisation of such approaches on mobile devices is a potentially promising research and engineering direction for the future.

TABLE III. Ranking of influencing factors on feature-based image-to-geometry registration procedures.

Order	*Equal radiometric conditions*	*Radiometric variances*
1	Input external orientation	Feature detection and description
2	Feature detection and description	Input external orientation
3	Sensor quality	Image processing
4	Internal parameterisation	Sensor quality
5	Camera calibration	Internal parameterisation
6		Camera calibration

The conducted performance measurement, in terms of computation speed and runtime in this paper, is of a preliminary nature for many reasons. First, each algorithm utilised in the interactive framework can be improved by GPU computing (such as OpenCL, CUDA on Tegra) to obtain faster computations. Second, the most severe computational bottleneck in the process is the large size of the underlying textured surface geometry. Runtimes can be improved by preprocessing the geometry more carefully, reducing the vertex count (possibly at the expense of some centimetres of geometrical accuracy), and separating radiometric (photograph texture) and geometric (vertex positions) information in the model, which is indicated by the results of this study. Furthermore, operational parameters of computational speed and energy consumption are possibly strongly correlated. This aspect, as well as further performance studies and improvements, are the subject of short-term future studies currently in progress in collaboration with other institutes which actively contribute to the research in this field.

User Intervention

The main contribution of this paper – the proposed interactive registration framework – offers various modes of user intervention and quality feedback to improve upon unreliable, previously obtained, fully automatic registration results. Careful user intervention allows accounting for unfavourable radiometric effects in image-to-geometry registration. This makes it possible to register challenging outdoor photographs to a level of accuracy that is comparable with diffusely illuminated photographs with simple imaging conditions. The comparison of the interpretation 3D mapping results in Fig. 12 support this conclusion. The framework can be extended in future with further user input, such as the demarcation of lit and unlit areas or the input of the view horizon, which simplifies the registration and resolves common view ambiguities (for example, the upwards view direction). Although techniques for the recalculation of natural lighting within computer graphics exist to

automatically resolve the illumination ambiguity of the DOM texture, this introduces uncertainty into the scene lighting approximation. Furthermore, it significantly adds to the DOM pre-processing procedure. Therefore, a user-guided, image-based illumination map may conceptually be simpler to realise.

Additional User Studies

Further field studies are necessary to assess the usability of the technique for repeated fieldwork at the same location. Additionally, an extended user study is important to quantify the acceptance of the presented approach within the target audience for everyday use. Furthermore, the methodology is applicable to other applications within the geosciences, such as glacial observations in climate studies, which needs a careful assessment of the limits of the given registration techniques under radiometrically challenging conditions from strongly reflecting surfaces. The mobile technology presented in this paper also allows an expansion of virtual fieldtrip ideas, as proposed by McCaffrey et al. (2010). The developed mobile tool can simply be distributed to fieldtrip participants to share data and interpretations in a consistent context. Moreover, apart from 2D interpretation mapping (such as geological sections), photograph registration also allows students to mark start-and-end points of 1D vertical measures (for example, sedimentary logs) with overview images. Then, by projecting the line segment spun by both markers onto the surface geometry, the related logs can be mapped in 3D in a sophisticated fashion. Incorporating and connecting visual logging applications (such as Strataledge) with the registration is a logical future extension.

The presented techniques and algorithms have already been employed on three additional case studies. Given the collective experiences from these case studies, there are further practical factors that influence the feature-based image-to-geometry registration and the interpretation mapping. The registration is commonly applied to textured triangular-mesh geometry which is collected for a specific application domain (such as digital outcrop analogue studies within petroleum geology). The geometry is externally prepared for the scenario, which leads to cut-out parts of the surface model that are not of interest for the application domain but which are supportive, or indeed vital, for a successful registration. The lack of geometric information leads to a visual cluster of features in the rendered image, and some image parts without any features. As discussed by Gauglitz et al. (2011a), a homogeneous sample of features across the image potentially increases the correlation accuracy of algorithms, such as the one applied in the presented framework. This feature sample on the image cannot be controlled where actual content (namely, the textured surface) is missing. Furthermore, the quality of input signals, such as GNSS for positioning or the visual quality of the surface's texture, has a great influence on feature correlation and photograph registration. Fig. 13 illustrates this: although the initial pose approximation is good enough for the suggested process, the poor texture quality leads to erroneous feature correlations with every existing feature-matching method. Dynamic camera distortions, together with the lack of imaging parameter control (such as white balance or colour mapping) in many mobile cameras, are further influencing factors for accurate image registration with respect to photograph-acquisition modalities. Improved camera manufacturing and programmatic control of the imaging parameters for mobile devices may resolve such problems with future devices.

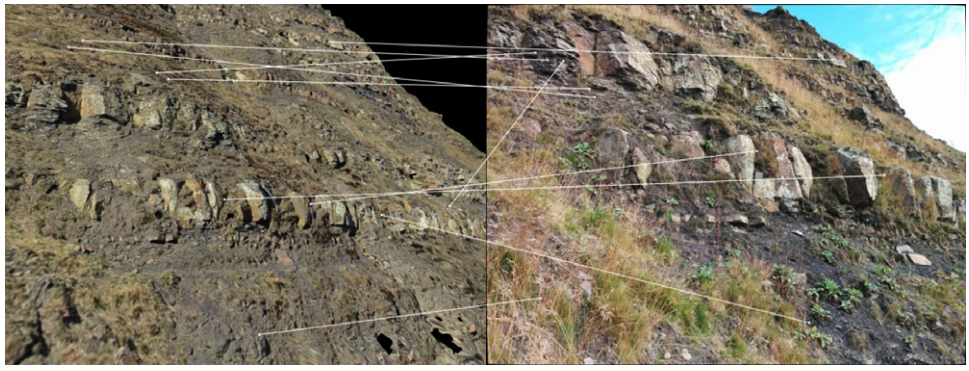

FIG. 13. Example of the negative influence of radiometric variance on the image-to-geometry registration procedure. Although, to the human eye, these two images look reasonably similar, feature-based methods fail – without exception – to produce acceptable registration results. Interpretations on this image would be without 3D value.

CONCLUSION

In this paper, a new framework for image-to-geometry registration on mobile devices has been introduced, specifically targeted at semi-automatic registration of outdoor images affected by challenging, changing environmental conditions. The framework accounts for common problems with real-world photograph registration onto pre-acquired textured surface models, such as DOMs, under uncontrollable illumination and radiometric conditions (previously referred to as *radiometric variance*).

The framework uses DEMs for assisted repositioning to adapt erroneous GNSS positioning of internal mobile-device sensors, which are shown to significantly influence the final registration accuracy. In order to account for the radiometric variance, several image-processing procedures and registration parameters are provided to the user in order to improve keypoint correlations. Additionally, the presented work improves the communication of the registration results to the user. Icon indicators and a five-stage coloured quality classification are employed to show the improvements between successive registration trials, as well as an overall measure for the final registration accuracy.

The technical results are assessed on a geological case study of Mam Tor, a turbidite outcrop in the Peak District, UK. There, the stratigraphic features have been interpreted on two image sets that were acquired in different field campaigns and under varying photograph-acquisition conditions. The results show a large improvement in the overall registration accuracy and success rate when employing the interactive registration framework, in comparison with static geolocation and fully automatic image-to-geometry registration.

Moreover, the paper contributes a detailed discussion on open research questions within outdoor image-to-geometry techniques, as well as the drawbacks and technical limitations of mobile devices for image-based fieldwork. Mobile-device sensors provide a large amount of georeferenced information that can be used in geological and geospatial applications in the field, but their consumer-grade electronics provide limited accuracy for georeferencing. As illustrated by the presented approach, sophisticated techniques from image analysis, computer vision and photogrammetry are necessary to provide a correct, georeferenced link

between the visual mobile-device information and the real world. Further research is particularly necessary for the correct calibration of mobile-device cameras to increase the image-to-image and image-to-geometry registration accuracy.

ACKNOWLEDGEMENTS

This research is part of the VOM2MPS project (no. 234111/E30), funded by the Research Council of Norway (RCN) and the FORCE consortium through Petromaks 2 and SAFARI. Data are collected and provided in the framework of SAFARI (www.safaridb.com).

REFERENCES

BELLIAN, J. A., KERANS, C. and JENNETTE, D. C., 2005. Digital outcrop models: applications of terrestrial scanning lidar technology in stratigraphic modeling. *Journal of Sedimentary Research*, 75(2): 166–176.

BLUM, J. R., GREENCORN, D. G. and COOPERSTOCK, J. R., 2012. Smartphone sensor reliability for augmented reality applications. In *Mobile and Ubiquitous Systems: Computing, Networking and Services* (Eds. K. Zheng, M. Li and H. Jiang). Springer, Berlin, Germany. 288 pages: 127–138.

BRADSKI, G. and KAEHLER, A., 2008. *Learning OpenCV: Computer Vision with the OpenCV Library*. O'Reilly Media, Sebastopol, California, USA. 580 pages.

BUCKLEY, S. J., ENGE, H. D., CARLSSON, C. and HOWELL, J. A., 2010. Terrestrial laser scanning for use in virtual outcrop geology. *Photogrammetric Record*, 25(131): 225–239.

CALLIERI, M., CIGNONI, P., GANOVELLI, F., MONTANI, C., PINGI, P. and SCOPIGNO, R., 2003. VCLab's tools for 3D range data processing. *4th International Symposium on Virtual Reality, Archaeology and Intelligent Cultural Heritage (VAST)*, 3: 5–7. 10 pages.

CORSINI, M., DELLEPIANE, M., PONCHIO, F. and SCOPIGNO, R., 2009. Image-to-geometry registration: a mutual information method exploiting illumination-related geometric properties. *Computer Graphics Forum*, 28(7): 1755–1764.

CORSINI, M., DELLEPIANE, M., GANOVELLI, F., GHERARDI, R., FUSIELLO, A. and SCOPIGNO, R., 2013. Fully automatic registration of image sets on approximate geometry. *International Journal of Computer Vision*, 102(1–3): 91–111.

DELLEPIANE, M., MARROQUIM, R., CALLIERI, M., CIGNONI, P. and SCOPIGNO, R., 2012. Flow-based local optimization for image-to-geometry projection. *IEEE Transactions on Visualization and Computer Graphics*, 18(3): 463–474.

DEY, S. and GHOSH, P., 2008. GRDM – a digital field-mapping tool for management and analysis of field geological data. *Computers & Geosciences*, 34(5): 464–478.

ENDEEPER, 2016. *Strataledge* – mobile system for rock core descriptions. http://www.endeeper.com/product/strata ledge [Accessed: 16th August 2017].

ENGE, H. D., BUCKLEY, S. J., ROTEVATN, A. and HOWELL, J. A., 2007. From outcrop to reservoir simulation model: workflow and procedures. *Geosphere*, 3(6): 469–490.

FERSTER, C. J. and COOPS, N. C., 2013. A review of earth observation using mobile personal communication devices. *Computers & Geosciences*, 51: 339–349.

FORSSÉN, P.-E., 2007. Maximally stable colour regions for recognition and matching. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 1–8.

GAUGLITZ, S., FOSCHINI, L., TURK, M. and HÖLLERER, T., 2011a. Efficiently selecting spatially distributed keypoints for visual tracking. *18th IEEE International Conference on Image Processing*, 1869–1872.

GAUGLITZ, S., HÖLLERER, T. and TURK, M., 2011b. Evaluation of interest point detectors and feature descriptors for visual tracking. *International Journal of Computer Vision*, 94(3): 335–360.

GAUGLITZ, S., SWEENY, C., VENTURA, J., TURK, M. and HÖLLERER, T., 2014. Model estimation and selection towards unconstrained real-time tracking and mapping. *IEEE Transactions on Visualization and Computer Graphics*, 20(6): 825–838.

GRAIL, 2017. *Photo Tourism – Exploring Photo Collections in 3D*. Graphics and Imaging Laboratory (GRAIL), Dept. of Computer Science and Engineering, University of Washington, Seattle, Washington, USA. http://phototour.cs.washington.edu/ [Accessed: 16th August 2017].

GUISLAIN, M., DIGNE, J., CHAINE, R. and MONIER, G., 2017. Fine scale image registration in large-scale urban LIDAR point sets. *Computer Vision and Image Understanding*, 157: 90–102.

HODGETTS, D., 2013. Laser scanning and digital outcrop geology in the petroleum industry: a review. *Marine and Petroleum Geology*, 46: 335–354.

HOWELL, J. A., MARTINIUS, A. W. and GOOD, T. R., 2014. The application of outcrop analogues in geological modeling: a review, present status and future outlook. *The Geological Society, London, Special Publications*, 387: 1–25.

JAZAYERI, I. and FRASER, C. S., 2010. Interest operators for feature-based matching in close range photogrammetry. *Photogrammetric Record*, 25(129): 24–41.

KEHL, C., BUCKLEY, S. J. and HOWELL, J. A., 2015. Image-to-geometry registration on mobile devices – an algorithmic assessment. *Proceedings of 3D NordOst*, Berlin, Germany. Pages 17–26.

KEHL, C., BUCKLEY, S. J., GAWTHORPE, R. L., VIOLA, I. and HOWELL, J. A., 2016. Direct image-to-geometry registration using mobile sensor data. *ISPRS Annals of Photogrammetry, Remote Sensing & Spatial Information Sciences*, 3(2): 121–128.

KEHL, C., BUCKLEY, S. J., VISEUR, S., GAWTHORPE, R. L. and HOWELL, J. A., 2017. Automatic illumination-invariant image-to-geometry registration in outdoor environments. *Photogrammetric Record*, 32(158): 93–118.

KRÖHNERT, M., 2016. Automatic waterline extraction from smartphone images. *International Archives of Photogrammetry, Remote Sensing and Spatial Information Sciences*, 41(B5): 857–863.

LEPETIT, V., MORENO-NOGUER, F. and FUA, P., 2009. EP*n*P: an accurate O(*n*) solution to the P*n*P problem. *International Journal of Computer Vision*, 81(2): 155–166.

LOWE, D. G., 2004. Distinctive image features from scale-invariant keypoints. *International Journal of Computer Vision*, 60(2): 91–110.

MAES, F., COLLIGNON, A., VANDERMEULEN, D., MARCHAL, G. and SUETENS, P., 1997. Multimodality image registration by maximization of mutual information. *IEEE Transactions on Medical Imaging*, 16(2): 187–198.

MCCAFFREY, K. J. W., JONES, R. R., HOLDSWORTH, R. E., WILSON, R. W., CLEGG, P., IMBER, J., HOLLIMAN, N. and TRINKS, I., 2005. Unlocking the spatial dimension: digital technologies and the future of geoscience fieldwork. *Journal of the Geological Society*, 162(6): 927–938.

MCCAFFREY, K. J. W., HODGETTS, D., HOWELL, J., HUNT, D., IMBER, J., JONES, R. R., TOMASSO, M., THURMOND, J. and VISEUR, S., 2010. Virtual fieldtrips for petroleum geoscientists. *Geological Society, London, Petroleum Geology Conference Series*, 7: 19–26.

MIDLAND VALLEY, 2017. *Digital Field Mapping*. http://www.mve.com/digital-mapping [Accessed: 16th August 2017].

MIKOLAJCZYK, K. F., TUYTELAARS, T., SCHMID, C., ZISSERMAN, A., MATAS, J., SCHAFFALITZKY, F., KADIR, T. and VAN GOOL, L., 2005. A comparison of affine region descriptors. *International Journal of Computer Vision*, 65(1): 43–72.

MOFFITT, F. H. and MIKHAIL, E. M., 1980. *Photogrammetry*. Third edition. Harper & Row, New York, USA. 648 pages.

MULLINS, J. R., HOWELL, J., KEHL, C. and BUCKLEY, S., 2016. Virtual outcrop models to multiple point statistics: improved reservoir modeling from virtual outcrops supported by digital field computing. *AAPG Annual Convention and Exhibition*, Calgary, Canada.

OLARIU, M. I., AIKEN, C. L. V., BHATTACHARYA, J. P. and XU, X., 2011. Interpretation of channelized architecture using three-dimensional photo real models, Pennsylvanian deep-water deposits at Big Rock Quarry, Arkansas. *Marine and Petroleum Geology*, 28(6): 1157–1170.

PINTUS, R. and GOBBETTI, E., 2015. A fast and robust framework for semiautomatic and automatic registration of photographs to 3D geometry. *ACM Journal on Computing and Cultural Heritage*, 7(4): article no. 23.

POWELL, M. J. D., 2006. The NEWUOA software for unconstrained optimization without derivatives. In *Large-scale Nonlinear Optimization* (Eds. G. Pillo and M. Roma). Springer, Berlin, Germany. 298 pages: 255–297.

PRINGLE, J. K., HOWELL, J. A., HODGETTS, D., WESTERMAN, A. R. and HODGSON, D. M., 2006. Virtual outcrop models of petroleum reservoir analogues: a review of the current state-of-the-art. *First Break*, 24(3): 33–42.

QUAN, L., 2010. *Image-based Modeling*. Springer, New York, USA. 251 pages.

RARITY, F., VAN LANEN, X. M. T., HODGETTS, D., GAWTHORPE, R. L., WILSON, P., FABUEL-PEREZ, I. and REDFERN, J., 2013. LiDAR-based digital outcrops for sedimentological analysis: workflows and techniques. *Geological Society, London, Special Publications*, 387: 153–183.

REMONDINO, F., RIZZI, A., BARAZZETTI, L., SCAIONO, M., FASSI, F., BRUMANA, R. and PELAGOTTI, A., 2011. Review of geometric and radiometric analyses of paintings. *Photogrammetric Record*, 26(136): 439–461.

RICHET, R. B. J., ADAMS, E. W., MASSE, J.-P. and VISEUR, S., 2011. Numerical outcrop geology applied to stratigraphical modeling of ancient carbonate platforms: the Lower Cretaceous Vercors carbonate platform

(SE France). *Outcrops Revitalized, Tools, Techniques and Applications (SEPM Special Publication)*, 10: 195–210.

ROSTEN, E., PORTER, R. and DRUMMOND, T., 2010. Faster and better: a machine learning approach to corner detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 32(1): 105–119.

SCHILLING, A., COORS, V. and LAAKSO, K., 2005. Dynamic 3D maps for mobile tourism applications. Chapter 15 in *Map-based Mobile Services* (Eds. L. Meng, T. Reichenbacher and A. Zipf). Springer, Berlin, Germany. 260 pages: 227–239.

SIMA, A. A. and BUCKLEY, S. J., 2013. Optimizing SIFT for matching short wave infrared and visible wavelength images. *Remote Sensing*, 5(5): 2037–2056.

SIMA, A. A., BONAVENTURA, X., FEIXAS, M., SBERT, M., HOWELL, J. A., VIOLA, I. and BUCKLEY, S. J., 2013. Computer-aided image geometry analysis and subset selection for optimizing texture quality in photorealistic models. *Computers & Geosciences*, 52: 281–291.

SNAVELY, N., SEITZ, S. M. and SZELISKI, R., 2006. Photo tourism: exploring photo collections in 3D. *ACM Transaction on Graphics (TOG)*, 25(3): 835–846.

SOTTILE, M., DELLEPIANE, M., CIGNONI, P. and SCOPIGNO, R., 2010. Mutual correspondences: a hybrid method for image-to-geometry registration. *Eurographics Italian Chapter Conference*: 81–88.

SOUTHERN, S. J., MOUNTNEY, N. P. and PRINGLE, J. K., 2014. The Carboniferous southern Pennine basin. *UK. Geology Today*, 30(2): 71–78.

SWEENY, C., FLYNN, J., NUERNBERGER, B., TURK, M. and HÖLLERER, T., 2015. Efficient computation of absolute pose for gravity-aware augmented reality. *IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*, Adelaide, Australia. Pages 19–24.

TORR, P. H. S. and ZISSERMAN, A., 2000. MLESAC: a new robust estimator with application to estimating image geometry. *Computer Vision and Image Understanding*, 78(1): 138–156.

VISEUR, S., ROUDAUT, R., BERTOZZI, R., CASTELLI, M. and MARI, J.-L., 2014. 3D interactive geological interpretations on digital outcrops using a touch pad. *1st Vertical Geology Conference (VGC)*, Lausanne, Switzerland. Pages 109–113.

WALLIS, K. F., 1974. Seasonal adjustment and relations between variables. *Journal of the American Statistical Association*, 69(345): 18–31.

WALTHAM, A. C. and DIXON, N., 2000. Movement of the Mam Tor landslide, Derbyshire, UK. *Quarterly Journal of Engineering Geology and Hydrogeology*, 33(2): 105–123.

WANG, J., SCHINDLER, G. and ESSA, I., 2012. Orientation-aware scene understanding for mobile cameras. *ACM Conference on Ubiquitous Computing*: 260–269.

WOLNIEWICZ, P., 2014. SedMob: a mobile application for creating sedimentary logs in the field. *Computers & Geosciences*, 66: 211–218.

Résumé

Beaucoup d'applications en géosciences requièrent de plaquer des photos sur des surfaces. Cet article présente une méthode de plaquage interactive effectuée sur tablettes, c'est-à-dire sous des conditions de performance informatique limitée et d'acquisition d'images difficile. Elle utilise conjointement des Modèles Numériques de Terrain (MNT) et des données d'orientations et positions issues des capteurs de la tablette. Des heuristiques de calcul fondées sur les connaissances actuelles en recalage d'images retournent des paramètres permettant d'évaluer les résultats. Cette approche a été appliquée sur deux jeux d'images d'un même site, prises à deux moments différents. Les interprétations des images sont plaquées directement sur la surface topographique et leur précision est évaluée. Ces tests montrent une amélioration significative du calage des photos et de la précision du positionnement des interprétations sur la surface topographique. Le côté interactif de cette approche semi-automatisée permet d'obtenir des résultats nettement supérieurs à ceux de méthodes comparables entièrement automatisées.

Zusammenfassung

Die Projektion von Fotos auf Oberflächengeometrie ist für viele Anwendungen der Geowissenschaften von großer Bedeutung. Methoden zur Berechnung dieser 2D – 3D Transformation wurden bereits ausführlich

untersucht, unter besonderer Berücksichtigung von problematischen Abbildungseinflüssen und ungenauer, initialer Lokalisierung. Dieser Artikel präsentiert ein interaktives Framework zur direkten, punktebasierten Projektion von Bildinformationen auf Oberflächenmodelle in 3D auf mobilen Endgeräten. Problematische Abbildungseinflüsse der realen Naturfotographie sowie die begrenzte Rechenkapazität mobiler Endgeräte sind dabei von zentraler Bedeutung. Das präsentierte Framework benutzt frei-verfügbare, digitale Geländemodelle (dt.: DGM; eng.: DEMs), Sensordaten zur Positions- und Orientierungsbestimmung mobiler Endgeräte und heuristische Berechnungen zur automatischen Evaluierung, Resultatwiedergabe und dem visuellem Feedback. Der hierbei präsentierte Ansatz führt bestehendes Wissen und neue Methoden zur Bild-zu-Geometrieregistrierung zusammen. Das entwickelte Framework ist auf zwei zeitlich-verschiedenen Bilddatensätzen der "Mam Tor"-Gebirgsformation in Derbyshire (UK) angewandt und evaluiert worden. Die bildbasierten, geologischen Interpretationen werden hierbei auf ein texturiertes Oberflächenmodell (3D) projeziert und deren Projektionsgenauigkeit qualitativ gemessen und beurteilt. Die Experimente zeigen, dass die Anwendung des interaktiven Frameworks zu erkennbaren Verbesserung der Genauigkeit der Oberflächenregistrierung der Bilder sowei deren Interpretationen führt. Der halbautomatische, Anwender-gesteuerte, interaktive Ansatz ist vergleichbaren, vollautomatischen Bildregistrierungsansätzen, unter statistischen Genauigkeitskriterien, messbar überlegen.

Resumen

La correpondencia de fotografías a la geometría de la superficie es un procedimiento importante en muchas aplicaciones dentro de las geociencias. Este artículo propone un marco de trabajo interactivo para la correspondencia de imagen a geometría basada en características que funciona directamente en dispositivos móviles, en condiciones de imágenes difíciles y con hardware de rendimiento limitado. Se utilizan modelos de elevación digital (DEM) disponibles junto con datos de sensores de posición y orientación móviles. Integra cálculos heurísticos sintetizando los conocimientos disponibles en la literatura de registro actual para la evaluación de resultados y la retroalimentación. El método se evalúa en dos conjuntos de datos de imágenes captados en ocasiones separadas. Sus interpretaciones se asignan a un modelo de superficie texturizada lídar y la precisión de la proyección se evalúa cualitativamente. Los experimentos muestran una mejora significativa de precisión en los resultados de registro de fotografía, así como la fiel correspondencia de la imagen a la geometría de la superficie subyacente. Este enfoque interactivo semi-automático, guiado por el usuario es superior a otros métodos comparables de correspondencia totalmente automáticos.

摘要

在几何表面模型上映射影像纹理是地球科学中许多应用的重要步骤。本文提出了一种基于特征的影像到几何模型纹理映射的交互式方法，可以在困难成像条件及有限的硬件性能情况下，直接应用于移动装置。该方法利用可以公开获得的数字高程模型 (DEM) 以及行动装置中的位置和方位传感器数据, 使用启发式计算整合结果评估与反馈, 综合现有文献中之既有知识。本方法以两组在不同情况获得的影像进行评估, 将影像纹理映射到一个以激光扫描获取的具有纹理的表面模型，并进行定性评估映射精度。实验该方法对影像的配准精度有显著提高，　能够将影像纹理准确地映射到表面几何模型上去。这种半自动、使用者主导、交互式的方法, 优于类似情况下的全自动配准方法。