



Australian Government

Department of Defence

Defence Science and
Technology Organisation

Registration for Tactical Imagery: An Updated Taxonomy

R. Bruce Backman

**Weapons Systems Division
Defence Science and Technology Organisation**

DSTO-TR-1855

ABSTRACT

Registration is central to Image Processing problems which use Tactical Imagery. Any application, which involves comparing two or more images, requires some type of Registration algorithm. These algorithms have evolved over the years and are generally grouped into three categories: 2D, 3D, and the reasonably unusual combination of 2D/3D. An updated classification (or taxonomy) for the diverse collection of algorithms is presented here and is described in detail. Also, two new algorithms are elucidated: the terrain cube and the Hybrid registration method.

Many examples are given demonstrating the usefulness of this taxonomy and algorithms. The Medical Imaging field is the source for many of these examples, as numerous algorithms have their origin there. Complementary Military Imaging examples are also presented and described in terms of relevant platforms.

RELEASE LIMITATION

Approved for public release

Published by

*Weapons Systems Division
DSTO Defence Science and Technology Organisation
PO Box 1500
Edinburgh South Australia 5111 Australia*

*Telephone: (08) 8259 5555
Fax: (08) 8259 6567*

*© Commonwealth of Australia 2006
AR-013-641
April 2006*

APPROVED FOR PUBLIC RELEASE

Registration for Tactical Imagery: An Updated Taxonomy

Executive Summary

Many Defence applications require imaging or images in some form. It is prudent to understand the role played by Image Registration. This is the branch of Image Processing, and to some extent Computer Vision, focussed on the issue of aligning images. This alignment is necessary due to the movement of the target object, to the imaging platform (missile or aircraft, typically), and the presence of clutter such as clouds, thermals, and countermeasures. Good registration is critical for tactical applications such as superresolution, targeting, surveillance, and intelligence gathering. Additionally, the generation of Digital Point Positioning Data Bases (DPPDBs) for targeting by, for example, JASSM air-to-ground missiles, requires highly accurate, registered imagery.

Understanding what makes good registration requires an overview, or taxonomy, of the types of registration problems typically encountered in defence.

This updated taxonomy is constructed based on the registration needs from tactical imagery. The author believes that the origins of these images are sufficiently distinct from commercial, intelligence, and strategic imaging, to make a unique approach desired. Key highlights include:

1. Description of tactical imaging domain
2. Extensive definition of registration
3. Taxonomy of registration algorithms
 - a. 2D Image Matching (Multimodal/Template/Viewpoint/Temporal/Hybrid)
 - b. 3D Image Registration
 - c. 2D/3D Combined Registration
 - d. Video Registration

This taxonomy may assist in defining the scope of contracts and in understanding the technical challenges to be overcome in imaging projects. Additionally, collecting algorithms written in disparate areas of defence, both by in-house and contractor personnel, based on these groupings may help with the re-use of software modules.

Two new ideas, the terrain cube and the Hybrid class of registration algorithms, are introduced and described. The terrain cube offers a way of integrating 3D information which may come from many sources; such as DTED (Digital Terrain Elevation Data) and dynamic terrain maps generated on-the-fly from UAVs. This may be militarily useful in situations such as targeting moving vehicles in the air-to-ground scenario. The Hybrid algorithm is a way of potentially deriving highly accurate spatial information about moving ground targets from various sensors in the tactical environment. These include, for the purposes of defining the technique, a stationary reconnaissance vehicle and an overflying UAV. Further work on developing applications for this algorithm will be the author's next focus.

Authors

R. Bruce Backman

Weapons Systems Division

R Bruce Backman completed a B.Sc.(cum laude) in Mathematics and Computer Science from Mount St. Vincent University in Halifax, Canada. He then worked for Exxon/Mobil designing Computer Networks. Following that experience, he migrated to Australia and began PhD studies at The University of Western Australia in Perth, Australia. There he studied Computer Science and Surgery, supervised by Professor Fiona Wood, 2005 Australian of the Year. During this time he published a number of papers and completing a thesis entitled "A Temporal 3D-Registration Framework for Computer-Integrated Surgery". The focus of the thesis was to provide scientific evidence to show that early surgical intervention in severely burn-injured patients improves their mortality rates. He started with DSTO in 2003 and was assigned to the Electro-Optical Seekers group in the Weapons Systems Division. He has interests in Unmanned Aerial Vehicles (UAVs), Image and Video Registration, Camera Calibration and Multi-Camera Systems.

Contents

1. INTRODUCTION	1
1.1 Overview	1
1.2 Motivation	1
1.3 Outline of the Paper	2
2. DEFINING 2D IMAGE REGISTRATION	2
2.1 Mathematical Basis	2
2.2 Computer Science Abstraction	4
3. A REGISTRATION TAXONOMY FOR TACTICAL IMAGERY	6
3.1 2D Image Registration	6
3.1.1 Multimodal	7
3.1.2 Template	8
3.1.3 Viewpoint	10
3.1.4 Temporal	11
3.1.5 Hybrid (Temporal/Viewpoint)	12
3.1.6 2D Registration Summary	15
TABLE 1. SUMMARY OF 2D IMAGE REGISTRATION TECHNIQUES WITH SUGGESTIONS FOR APPLICATIONS.....	15
3.2 3D Image Registration	16
3.3 2D/3D Combined Image Registration.....	18
3.3.1 General	18
3.3.2 GeoLocation	18
3.4 Video Registration.....	21
3.4.1 Video to Reference Imagery or 3D models: GeoLocation.....	21
3.4.2 Video to video registration.....	21

3.4.3	Frame to frame registration.....	22
4.	CONCLUSIONS	22
5.	ACKNOWLEDGEMENTS.....	23
6.	REFERENCES.....	23

1. Introduction

1.1 Overview

Registration, the discovery of the transformation from one image to another, is a central problem that must be addressed in many areas of tactical imaging¹. Knowing this transformation allows us to dynamically adjust the performance of real-time seekers and to plan for the interpretation of *metadata*² sent concurrent with image and video streams. Another critical application is the need to enhance the resolution of a still image extracted from a video sequence or of the video sequence itself. This enhanced resolution is possible because the spatial correlations between successive image frames can be exploited. Such a multi-frame reconstruction process is usually called a superresolution (SR) reconstruction. The study of the performance of registration algorithms is difficult as allocated by E. Cuchet³ [1], an expert in the field: “It is very difficult to study registration algorithms on real data, as the ideal transformation to be found is unknown”.

1.2 Motivation

Registration problems are encountered in many tactical military applications: missile seeker systems, mission planning, and UAV imagery, to name a few examples [2]. Registration problems in these applications (as distinct from the Surveillance or Intelligence domains) are typically complicated by the presence of poor quality, small-sized, brief image sequences. This is due to the highly dynamic movements encountered by the sensors. Also, most image acquisition is complex in nature, often obscured by counter-measures, smoke, fog, and other environmental factors. Most well-known

¹ The domain of tactical imaging is defined to include those images and image sequences generated from weapons and equipment used in military manoeuvres. For example, the high rate transmission of low-quality, infrared images from an air-to-air missile is considered tactical imagery. Higher-resolution imagery derived over a longer period of time with much higher accuracy is defined to be strategic imagery. This may include satellite imagery used for intelligence purposes.

² Metadata is the term used to refer to non-image data that is transmitted in parallel with image or video sequences. This data may contain much information about the orientation, calibration, and power of the image sensor and, in some UAVs and aircraft, information about aerodynamic status. Metadata is commonly encapsulated in the video stream using techniques such as Video Blanking Intervals. This allows the metadata and images to be transmitted on the same communication channel.

³ E. Cuchet is a world-renowned researcher in medical applications of image registration; specifically those for neurosurgery. She is based at INRIA (Institut National de Recherche en Informatique et en Automatique) in France.

registration algorithms expect well-defined easily identifiable registration features. In many cases, these difficult imaging conditions are encountered, causing even heavily tested algorithms to falter. One example of this is the case where *mirages*, or highly reflective water-like regions, are encountered on low-level, high-speed flights during high heat conditions in desert-like regions such as the Australian Outback. This causes difficulty for registering images from both infrared (IR) and electro-optical (EO) sensors. These sudden high-contrast and relatively large regions tend to quickly dominate or *washout* the targets being tracked.

1.3 Outline of the Paper

Now that the introduction and motivation have been established we proceed to the second major section of the paper: Defining 2D Image Registration. In this section, as a way of describing the problem domain, some simple, well-known mathematical descriptions of the problem are given. And, most importantly, the author characterises the registration problem using an application of *Abstraction* taken from Theoretical Computer Science. It is intended that this abstraction serve as a way of conceptualising the registration process as well as helping to focus on the improvement of algorithms in the tactical imagery domain.

The third major section of the paper: A Registration Taxonomy⁴ for Tactical Imagery provides a directory for classifying registration problems (both 2D and 3D). The specific example of GeoLocation is described in order to show how the taxonomy can be applied to current tactical problems such those encountered with the use of small UAVs. The Conclusions form the fourth major and final section of the paper.

2. Defining 2D Image Registration

2.1 Mathematical Basis

A simple way of defining 2D rigid-body image registration is to define a transformation, or mapping, that maps points (or pixels) in one image to their corresponding point in the other. This mapping is accomplished both spatially and with respect to intensity. Modersitzki⁵ [3] demonstrates this with a generalised approach, given here. We define a distance measure $D: \text{Img}(d)^2 \rightarrow \mathbb{R}$ and two images $R, T \in \text{Img}(d)$ (typically referred to as the Reference and Template images). Two mappings (spatial and intensity⁶) are given:

⁴ The study of the general principles of scientific classification.

⁵ In his book, "Numerical Methods for Image Registration", Oxford University Press, 2004.

⁶ The term *spatial* here is used to refer to the pixel-by-pixel mapping from one image (Reference) to a second image (Template). If the two images are from different sources the *intensity* values will need

$\varphi: \mathbb{R}^d \rightarrow \mathbb{R}^d$ and $g: \mathbb{R} \rightarrow \mathbb{R}$ such that $D(R, g \circ T \circ \varphi)$ is minimised. To improve this process and add more *depth* or texture to the images, feature spaces must be populated in order to accomplish a better-defined registration. This can be described as follows: we have $m \in \mathbb{N}$ and the features $F(R, j)$ and $F(T, j)$, $j = 1, \dots, m$. A transformation $\varphi: \mathbb{R}^d \rightarrow \mathbb{R}^d$ is defined such that $F(R, j) = \varphi(F(T, j))$, $j = 1, \dots, m$. This then leads us to the definition of a *distance measure*:

$$DLM[\varphi] := \sum_{j=1}^m \left\| F(R, j) - \varphi(F(T, j)) \right\|_f^2$$

where $\|\cdot\|_f$ denotes a norm on the feature space, e.g., $\|\cdot\|_f = \|\cdot\|_{\mathbb{R}^d}$, if the features are locations of points. The selection of these feature points in the images is critical; giving rise to many algorithms for detecting corners, edges, lines, and other geometric measures.

However, practical implementations of these feature spaces are often difficult given the sparse nature of many tactical image sequences. Generally, we employ the affine transformation, from the projective transformation superset, to model the changes between two images. In our use of the affine transformation we assume that any change between images is *rigid body*⁷ [4] and not the much more difficult elastic problems often found in the medical imaging domain. For example, in the *anatomical-atlas* problem where a *normal* MRI (Magnetic Resonance Imaging) image of a brain is deformed to match, region by region, to that of the pathological, or diseased, sample.

An affine transformation is any transformation that preserves collinearity (i.e., all points lying on a line initially still lie on a line after transformation) and ratios of distances (e.g., the midpoint of a line segment remains the midpoint after transformation) [5]. Using MATLAB® to solve these registration problems often involves extensive use of the *tdata*⁸ data structure used by the Image Processing Toolbox. This data structure includes all attributes of the affine transformation as data members making the passing of registration parameters and re-use of existing registration algorithms more efficient.

to be mapped as well as the pixel values. This mapping (intensity) does not happen routinely in tactical applications.

⁷ Transformations which align volumes and surfaces are considered to be rigid body ones if they are constructed by assuming the movement of selected points on the body will represent the path that all points on that body will follow. This path must be fully described by a translation and subsequent rotation represented by an orthonormal matrix of determinant one.

⁸ The *tdata* data structure is used extensively within MATLAB when dealing with affine, and other, transformations. The *maketform* function generates the data structure which can then be used as input into other programs such as *imtransform*.

To theoretically consolidate the great numbers of algorithms for registration problems we now provide an abstraction of the problem which may assist implementation and integration of disparate algorithms.

2.2 Computer Science Abstraction

Now that the mathematical representation of image registration has been described, we turn to a functional abstraction of image registration for the purpose of eventual implementation. It is well accepted that Abstraction is the central tenet of Computer Science. Abstraction plays the critical role of determining how much detail is incorporated in the analysis of problems. A poorly abstracted design could be one where so much detail is incorporated that the *big picture* is lost. An equally poorly abstracted design retains little detail, which could be used for an implementation. The phrase *broad* abstraction applies to ways of observing complex, systemic patterns and logic. *Narrow* abstraction refers to the provision of close-up analytical assessments of well-defined problems with an emphasis on implementation.

One broad abstraction of image registration, as described in Figure 1, has been proposed by Ng and Ibanez in the recent text edited by Yoo [6]. It provides a framework with which we can understand how algorithms can be designed and tested for their efficiency. The goal of this abstraction is to describe the spatial mapping that brings the *Template Image* into alignment with the *Reference Image*.




Figure 1 – A functional abstraction for the registration process as described by L. Ng and L. Ibanez in the text “Insight into Images” edited by T. Yoo. Each processing block can be optimised for the type of imagery being used other tactical constraints such as high speed “burst” communications over a short time period.

To generate the first iteration of this registration abstraction, a *null* transform data structure and interpolation kernels are applied to the compare image. This allows the Root Mean Square⁹ (RMS), or other metric module, to return the baseline *worst-case* value. The RMS module is the most important one in the abstraction. It determines the success of optimisation and ultimately the entire logic flow. The optimisation module is then triggered to produce a new transform data structure, which is then matched to a compatible interpolation kernel and applied to the compare image. Ultimately, even a simple RMS¹⁰ (one consisting of individual mean-squared differences between pixel intensities) will generate a fitness value¹¹ acceptable enough to deny entry into the optimisation stage; effectively ending the registration process.

In Section 1.2 the problems with tactical imaging, for example, low-flying missiles, dramatic changes in scene composition, or rapidly vibrating tactical UAVs, were described [7]. In terms of this abstraction, we must add two additional modules—Template Image Pre-Processing and a Temporal Filter. These two modules work together to *improve* the quality of the Template Images. The rate of image acquisition in Tactical Systems allow us to *buffer*¹² or build up a small sequence (say five images) of template images before we apply the registration transformation and interpolation modules. Buffering and assessing the quality of compare images improve the application of registration algorithms using tactical imagery.

Each of the two images, reference and template, are typically generated using a sensor which produces a *continuous* signal. Generally, in the case of digital cameras, this signal is represented, or mapped, onto a *discrete* digital array. These digital arrays are then directly manipulated using computer image processing programs. However, during the registration process there is a subtle, often overlooked, sub-process which deals with the situation when, after manipulation by the Transform module, the template image is no

⁹ The measure of Registration Error is a complex issue. Isolating the true registration error component from other systemic influences is difficult. As a starting point, the grayscale values of the pixels are compared and a simple Euclidean Root Mean Square error is computed.

¹⁰ A simple RMS would be bounded between 0 and 255, the grayscale range of an 8-bit image. A value of 2.4 would indicate a close match. A value of 50 would be a coarser fit.

¹¹ The term fitness value comes from the Computer Science field of Genetic Algorithms and Neural Networks. It is a generic term which refers to an algorithms genetic *fitness* and therefore whether or not it should continue to be considered.

¹² Missiles, UAVs, aircraft and other Tactical sensor platforms move quickly and have rapidly changing scenes compares to Strategic imagery. This buffering allows the calculation of *keyframes*, those which provide the high quality information to the movement encountered in the image sequence. In the problem of 3D point-cloud reconstruction (terrain mapping) based on these images, these keyframes allow for the correction of *projective drift*; an algorithmic anomaly which distorts the calculation of 3D points from comparing 2D images.

longer a discrete-based image, having become continuous through the application of sub-pixel manipulations.

In order to estimate the success of the registration algorithm this discontinuity must be addressed. In many cases the discrete model is crude. This overshadows the effect of the registration often resulting in temporal artefacts. The higher the digital precision of the image model the less effect this mapping has on the registration. Generally the RMS module of the algorithm should reflect a confidence interval, or some other error metric, which describes the effect of this mapping.

We now present a detailed proposed taxonomy for registration problems using tactical imagery. The taxonomy is based on the ease with which algorithms within categories can be modified to minimise development time and enhance value from re-use of software modules.

3. A Registration Taxonomy for Tactical Imagery

The term Taxonomy is used in the biological sciences to describe the inherited relationships between organisms, which allow them to be grouped into smaller and smaller categories. In this application the term is used to describe the algorithmic relationships that different types of registration problems have with one another. The top four groupings in our registration taxonomy are: two-dimensional (2D), three-dimensional (3D), (two/three)-dimensional (2D/3D), and video registration algorithms.

Most other groupings of registration algorithms in the tactical imaging literature focus on variations of the two-dimensional problem. However, sensors (cameras) are being used in new tactical applications, such as UAVs and LADAR (Laser Detection And Ranging), which allow us direct access to clouds of three-dimensional points, which demand different algorithms for registration.

3.1 2D Image Registration

The author proposes a well-known taxonomy for 2D image registration first espoused by Lisa Brown in her seminal survey paper written in 1992 [8]. The first four descriptions are based on Brown's however they are explained here in a tactical context. The final category, Hybrid (Temporal/Viewpoint), has been proposed by the author and is considered to be unique to this problem domain.

1. Multimodal
2. Template
3. Viewpoint
4. Temporal
5. Hybrid (Temporal/Viewpoint).

A brief description of each category follows.

3.1.1 Multimodal

Multimodal image registration is defined as the registration of images of the same scene acquired from different sensors. This can be the case where a remotely sensed target area is imaged from a UAV or satellite using cameras operating in the infrared, ladar, visual, and radar modalities. The key to successful multimodality image registration is the acquisition of either *intrinsic* or *extrinsic* features.

Intrinsic features are those derived from the calculation of image-derived measures. These can consist of things such as edges, contours, or lines. Additionally, higher-level features such as statistically derived vectors can be computed. These may include the correlation of gray-level means, centroid and principal axes calculations, as well as higher order moments of distribution such as kurtosis and skewness. Structural features may also be translated into data structures such as graphs. These abstract structures can then be matched against previously constructed *models*. These may be geographic, object, or in the civilian case, anatomic [9, 10].




Figure 2 – Image acquired from Electro-Optical and Infrared cameras. The checkerboard pattern of the fuel storage tank (shown in red) was used as an intrinsic feature for registration. In the Infrared image the similarly sized object was isolated and used to match up the two images. Images were used courtesy of Penn State University's Computer Science Department.

Extrinsic features are those derived from elements of the image, which have been added *a priori*. In the case of UAV trials imagery features may consist of brightly coloured plastic sheets, which are laid on the ground during the aircraft overflights. During targeting runs, certain tags may be superimposed into the images by an expert viewer—one privy to intelligence of the geographic and structural area. These tags may, in the case of Figure 2, highlight the Fuel Storage Tank in the bottom-left of the image.

The benefit of using multimodal registration is the added value given to the merged image by the introduction of the second modality¹³. In the case of the example (see Figure 2) we have an optical image of a highway, fuel tank, and runway, taken from a low-altitude UAV, aircraft, or missile. Although the structural intrinsic characteristics in the image can be derived (such as the outline of the fuel tank, road, etc.) crucial information such as if a small vehicle moving along the highway may be missed. Adding the second infrared modality may show a vehicle moving at speed (generating a heat plume), which could then be targeted.

3.1.2 Template

Template registration consists of finding a match for a reference pattern in an image. In the case of targeting or surveillance imagery, this would involve the identification of well-known features such as airports or fuel depots. In the case of civilian (mostly medical) image registration, this usually takes the form of the *anatomical atlas* problem where a patient's MRI (Magnetic Resonance Imaging) or CT (Computed Tomography) scan is matched against the template, or well-accepted, normal model. A commonly discussed and militarily useful application of template registration is encountered in the application of *Superresolution* (see Figure 3). This technique is accomplished by taking a number of low-resolution, poorly resolved images and interpolated the missing pixels, removing noise and blur to produce a single high-resolution image.

¹³ Most algorithms for semi-automatic or fully automatic multimodality registration have their origins in the medical imaging domain where, typically, CT and MRI (Computed Tomography & Magnetic Resonance Imaging) are combined to show disease progression. CT is the more spatially coherent modality and acts as an *anchor* for the MRI which more accurately displays dynamic physiology.




Figure 3 – A sample superresolved infrared fighter silhouette image is shown (top right). The high-resolution (480x640) image (top left) was artificially sub-sampled, using randomly selected offset vectors, into 10 low-resolution (120x160) analogues (bottom two rows). These images are then aligned using a feature map generated from a normalised cross correlation registration. The 10 images are then combined into one composite low-resolution image and linear mapping is applied to resample the image to be the same as the input image. This interpolation step “stretches” the gray-scale pixel values across the empty space created by upsampling the image. The RMS error measuring the difference in pixel values was 9.58 on a scale where 0 indicates no difference and 255 indicates a completely black image compared with a completely white one. Note that the superresolved image has had its values inverted for easier viewing.

In the example clearly defined template features (the fighter) are used making the registration quite simple. A naïve algorithm such as basic cross-correlation can produce fast high-quality results. The challenge with template registration is threefold. First, the identification of the template in the base image may be difficult requiring extensive *man-in-the-loop* intervention. Second, clutter, such as clouds, reflections, and the presence of countermeasures, may make even a well-defined initial template difficult to match in successive images. Third, seeker noise such as dome-heating, reflections, ghosting and dead pixels have to be incorporated in the registration model. Clearly, a template image pre-processing module and a temporal filter or buffer would ameliorate these effects.




Figure 4 – Figure on the left shows a UAV (Aerosonde®) imaging a stationary Fuel Truck from one position. The image on the right shows the same UAV imaging from a second position. This demonstrates the classical viewpoint registration problem. Knowing the variables associated with each image capture allows the efficient application of stereo techniques for the derivation of depth and shape.

3.1.3 Viewpoint

Viewpoint registration addresses the registration of approximately the same image scene taken from different viewpoints or positions. There are mathematically beneficial reasons for doing this. This includes the recovery of depth or shape information using, for example, epipolar geometry to constrain the calculation of the intersection of the respective optical axes. In this classical Computer Vision technique, an epipolar plane is defined which includes the epipolar lines in each image plane. These lines bisect the image planes—separating it into a top and bottom. This bisection line is broken into a left and right segment. The Centre of Projection line for each image goes through this centre point terminating in the point of interest on the target. The resulting search space is greatly reduced, that is, points on the target not contained in the epipolar plane are ignored. Viewpoint registration also facilitates video registration which is becoming much more popular as UAVs and other tactical platforms are integrated into the Network Centric Warfare Environment (see Figure 4). In this application, consecutive images in a sequence may reveal scenes in which mobile targets can be annotated.

Recently, the definition of viewpoint registration has been modified to include slightly different views of overlapping scenes. This is due to the increased use of UAVs for *terrain mapping* and related techniques to provide richer 3D information. An application of this is the technique of Mosaicing (see Figure 5). This involves *stitching* together images of terrain

using well-defined intrinsic or extrinsic features. Taking into consideration the perspective projection differences with each image acquisition, we quickly find that a dense 2D terrain mapping problem becomes a problem of 3D registration. This is because overlaying the 2D images, each taken with slightly different camera variables imply that 3D geometry is considered before completing the mosaic. Practical viewpoint registration problems often involve slight shifts in the defined image scene (see Figure 6). Adjustments to align these scenes prior to registration must be made for accurate results. Classical Viewpoint problems are not usually encountered in tactical imaging as the imaging platform often moves rapidly relative to the target resulting in large discrepancies. Also, any movement in the scene between image acquisitions renders this technique highly inaccurate.




Figure 5 – An example of a dense mosaicing image demonstrating an application of viewpoint registration. The Regions of Interest give an approximation of the size of the component images. The dashed-line boxes show the reconstruction grid for the mosaicing algorithm. Several hundred images were registered using Optical Flow – a technique based on the calculation of flow fields that are derived from the partial derivatives of the image values near the image boundaries. This is an elaborate example of the use of higher order intrinsic features. Images were used courtesy of Penn State University's Computer Science Department.

3.1.4 Temporal

Temporal registration algorithms are employed to work with imagery that has been generated by a stationary, fixed camera with a constant focal length imaging the same scene over a period of time (see Figure 7). Changes in the scene are registered to the initial or reference image. During the construction of the appropriate algorithms care must be used to distinguish between *global* and *local* changes in the scene. Global changes are things such as changes in lighting, contrast, or focal length. Local changes are those that can be modelled independently of the overall image. Applying a consistent algorithm to an extended temporal sequence may be quite difficult when global changes occur. The algorithm must be aware of these changes and communicate these new parameters to the registration modules (which may, ideally, be truly polymorphic). Global and local changes

should be considered in all types of image registration; however, they are more critical as the scene remain unchanged.




Figure 6 – Another example of a tactical viewpoint registration problem. This gunsight sequence, showing the destruction of a Serbian T-80 Tank during the Kosovo conflict using what is thought to be a Maverick air-to-ground missile, demonstrates two key points. Firstly, even though it is from a video sequence, it is not a temporal registration problem (this is described in the next section). Secondly, real life examples show that the viewing of the target is not exactly accurate. The crosshairs on the gunsight shift as the aircraft moves. In addition, real-time calculation of depth and shape information is a challenge. Currently these calculations are done on an off-line post-processing basis. These images were sourced from GlobalSecurity.Org's Intelligence Imagery database

In tactical applications this may be applied in the case of an Electro-Optic sensor (camera) mounted on a surveillance mast being used from a reconnaissance vehicle. The scene may be a major supply thoroughfare where observing the movement of military assets is of key importance. A complementary strategic example would be the positioning of a satellite in the same position to view the activity around a nuclear plant or other high value asset. The automation, or semi-automation, of feature extraction and tracking algorithms is highly desired in these applications. In practice, however, examples of pure temporal registration problems are unusual. Hybrid (Temporal/Viewpoint) cases are typically encountered.

3.1.5 Hybrid (Temporal/Viewpoint)

It is the author's assertion that most tactical registration problems are a fast-moving combination of the two classical temporal and viewpoint registration problems. An

example would be the viewing of aircraft targets from the seeker of an air-to-air missile. In this case both the target and camera are moving rapidly. The convergence of these two approaches severely limits the benefits generally attributed to each. Having a rapid five-frame sequence where the target moves quickly within the image frame disallows (or makes very difficult) the calculation of shape and depth information, which is the primary benefit of solving viewpoint registration problems. Similarly, although the camera may be focussed on the target aircraft continuously, the camera-equipped aircraft or missile is moving rapidly, disallowing the use of traditional temporal registration.




Figure 7 - An excerpt from a brief study done by the author to analyse the movement of a shockwave (shown by the red dots) through air after a controlled explosion. The camera is rigidly fixed during the trial and clearly shows the change in the scene--a classic temporal registration problem. Shockwave imagery courtesy of John Waschl, Terminal Effects Group/Weapons Systems Division.

Solving these registration problems in real-time is challenging. Typically, intensity thresholding is widely used because of its relatively fast algorithmic speed. Deriving higher-order features and classifying them is often too complex in these applications.

One solution (see Figure 8). for providing the richest possible information from the image sequences, that is, 3D points corresponding to image features, could be the following *decoupling* of registration vectors. If we take look at the situation where a stationary reconnaissance vehicle is imaging a moving target (another vehicle) we could assign the

movement of the vehicle in the image frame as belonging to a vector $\mathbf{T_Reg}_n$, $n = 1 \dots m$ where m is the number of images in the sequence. If we then had another sensor, in the case of this example a low-flying UAV with a camera looking straight down at the moving target, we could populate a second vector $\mathbf{V_Reg}_a$, $a = 1 \dots b$, where b is the number of viewpoint images taken from the UAV, with the epipolar points (as discussed previously) representing the vehicle. Typically, there may be $m = 100$ and $b = 10$. Then, to generate 3D points representing the vehicle, the two images, $\mathbf{V_Reg}_a$ and $\mathbf{V_Reg}_{a+1}$ could be considered the template and reference images. The movement of the vehicle could then be estimated by calculating the simple ratio: $m/b = 10$ to get the number of temporal assessments required to *match* each pair of viewpoint assessments. The information derived from the temporal image sequence will vary in quality depending on factors such as distance from the vehicle, angle above the ground, and intrinsic camera optical parameters. However, it may prove to be a valuable approach given the current tactical situation where there are imaging assets available to work together.




Figure 8 – Possible application of Hybrid registration. Here two cameras, one mounted on a stationary reconnaissance vehicle the other on the belly of a low-flying UAV, provide enough information to possibly provide a three-dimensional characterisation of the moving truck. Ideally, the stationary camera will pan, or follow, the target from side to side providing the greatest possible angular differences. Then the UAV tracks the target in the orthogonal path; following it from above. It is thought that this will provide the best precision solution.

3.1.6 2D Registration Summary

These registration techniques are summarised below in Table 1. Potential applications are described as well as advantages and disadvantages. Further research by the author will focus on the development of the Hybrid approach with applications to Terrain Mapping and Air-to-Ground targeting.

Table 1. Summary of 2D Image registration techniques with suggestions for applications.

	Advantages	Disadvantages	Summary	Application
Multimodal	Extra data to enrich scene	Requires user intervention to provide “anchor points”	2 sensors required as well as anchor points	Targeting, SAR/EOS
Template	Increase resolution, removes noise	Weak in dynamic imaging, latency	Matching to a reference image	Air-to-air aimpoint, SuperResolution
Viewpoint	Depth and shape recovery	Performs best with two sensors	Same scene – different viewpoints	Air-to-ground targeting support to ADF AIR 5418, 3D model matching
Temporal	Quick, accurate, spots changes	Assumes fixed camera	Same point registration	Surveillance
Hybrid	Highly resolved and accurate tactical scene	Complex	Leverages multi-sensors & NCW	Targeting, moving target aimpoint, littoral ships, ADF AIR 5418

3.2 3D Image Registration

Many military requirements demand accurate three-dimensional data representing unique intrinsic ground features and extrinsic, artificial features such as buildings and vehicles [11]. Traditionally, manned aircraft with conventional passive electro-optical camera technologies have done these mappings. Two cameras mounted far apart on the airframe allow for the acquisition of true three-dimensional ground data using well-known mathematical and image processing techniques (the epipolar constraint, as previously discussed). Assuming that the accuracy of this method is sufficient for tactical purposes we are presented with numerous, often quite dense (spatially), *clouds* of three-dimensional points. These individual point sets, each from different sensors, imaging the same 3D object, are then registered.

One representative technique, Arun's¹⁴ [12] algorithm, demonstrates one approach for 3D registration. Here, a rotation is applied to the first point set, then a translation, followed finally by a noise vector.

$$\{\mathbf{p}_i\}, \{\mathbf{p}'_i\}, \mathbf{p}'_i = \mathbf{R}\mathbf{p}_i + \mathbf{T} + \mathbf{N}_i.$$

We then form an optimisation (minimisation) function:

$$\Sigma^2 = \sum_{i=1}^n \left\| \mathbf{p}_i - (\mathbf{R}\mathbf{p}_i + \mathbf{T}) \right\|^2.$$

To solve this function we define some intermediate variables:

$$\mathbf{q}_i \triangleq \mathbf{p}_i - \mathbf{p}, \mathbf{q}'_i \triangleq \mathbf{p}'_i - \mathbf{p}', \mathbf{p}' \triangleq \frac{1}{n} \sum_{i=1}^n \mathbf{p}'_i, \mathbf{p} \triangleq \frac{1}{n} \sum_{i=1}^n \mathbf{p}_i.$$

This gives us a new minimisation condition:

$$\Sigma^2 = \sum_{i=1}^n \left\| \mathbf{q}'_i - \mathbf{R}\mathbf{q}_i \right\|^2.$$

¹⁴ The author used Arun's algorithm in his PhD thesis to find the relationship between two clouds of 3D points representing fiducials on the human body. In his work at DSTO he has continued to develop this code to be capable of handling very large datasets typically terrain maps. This code can be found here:

http://www.csse.uwa.edu.au/pub/robvis/theses/BruceBackman/code_and_data/fiducials/

$\hat{\mathbf{R}}$ must be found to minimise Σ^2 . $\hat{\mathbf{T}}$ is derived from $\hat{\mathbf{T}} = \mathbf{p}' - \hat{\mathbf{R}}\mathbf{p}$. The intermediate variables are combined as follows:

$$\mathbf{H} \triangleq \sum_{i=1}^n \mathbf{q}_i \mathbf{q}_i'^T.$$

Singular Value Decomposition (SVD) is then used to find the optimal solution: $\text{SVD}(\mathbf{H}) = \mathbf{U}\mathbf{W}\mathbf{V}^T$. If $\text{Det}(\mathbf{V}\mathbf{U}^T) = 1$, then $\hat{\mathbf{R}} = \mathbf{V}\mathbf{U}^T$, otherwise the algorithm fails. Therefore, in conclusion, we have $\mathbf{p}'_i = \hat{\mathbf{R}}\mathbf{p}_i + \hat{\mathbf{T}}_i$.




Figure 9 – Sample application of 3D registration used to integrate several terrain maps into one. The two overlap areas are isolated by hand then the registration algorithm is applied to generate the transformation matrices. These are then used to normalise all the 3D points in the scene. Items of interest such as moving vehicles and temporary buildings could be added to highly precise DTED data. Note that this is a simplistic 2D end-on view of a more complex 3D scene.

This demonstrates one algorithm for solving the general problem. It assumes no noise and, more critically, no obscuration in viewing the 3D object. This is a challenge with fast moving, low flying aircraft such as UAVs. For example, if the target is an airfield control tower and three 3D point sets are acquired from the requisite UAV sensors, there may well be large disparities in viewing angle from the sensor to the target. Therefore, it may then be required to submit a well defined partial piece of the target from one offending sensor to be registered with one more comprehensive 3D set from the other. Pre-processing will

be required to *match* the subset given by the first sensor to the corresponding points in the second. Once accomplished, since we are assuming a rigid-body object, we are able to continue to apply the algorithm as intended. The result is a feature-rich *terrain cube* comprising of, typically, DTED (Digital Terrain Elevation Data) and UAV-derived terrain information. As shown in Figure 9, two overlap areas (shown here in 2D, from the side-on perspective), A and B, can be registered to create the larger map. This could allow field commanders and analysts to incorporate targets of interest and subtle changes in terrain into the existing DTED models, which by definition, are generally historical elevation maps. Tactically, this means that the positions of easily relocatable targets such as mobile rocket launchers can be *refreshed* relative to highly accurate three-dimensional information.

3.3 2D/3D Combined Image Registration

3.3.1 General

Targeting applications which use imagery from aircraft cameras or missile seekers often address the challenge of matching 2D images with corresponding scenes in 3D models. This typically involves finding features in the images and then *back projecting* them through a known projective transformation into similar features in the 3D model (see Figure 10). To accomplish this accurately the *intrinsic* camera calibration parameters must be known. These can include focal length and distortion model coefficients. In the example (Figure 10), we see that edges (intrinsic features) are derived from the 2D image and these are correlated with the expected matching features in the 3D model to provide a level of confidence on the matching process.

3.3.2 GeoLocation

GeoLocation, or also geo-registration, is a specific case of 2D/3D registration typically encountered by UAVs and reconnaissance aircraft. An example of an algorithm for this type of problem is given by Shekhar & Chellappa [13]. They describe a relationship which, when given an image point, isolates a target point in the 3D world, typically a stationary (although not necessarily so) position. Given *metadata* from the host aircraft (in the form of measurements of the camera coordinate system, which includes the gimbal orientation) a homogeneous projection matrix can be formed which maps the 3D world points to image points. The information required for the formation of this matrix is:

Platform and gimbal position in latitude, longitude, and altitude (x_p, y_p, z_p)

platform orientation given by roll, pitch, and yaw ($\theta_r, \theta_p, \theta_y$)

gimbal orientation given by azimuth, elevation, and twist (ϕ_a, ϕ_e, ϕ_t)

camera internal parameters given by the horizontal and vertical fields of view (F_h, F_v) and image size in pixels (n_c, n_r). Note that no skew is assumed and that the principal point is the in the image centre.




Figure 10 - Intrinsic features, in this case, edges, are use to develop a map of an area of interest (Pre-processed Image). This map is then correlated to 2D projections from a known 3D projected model (Projected Model/Reference Image). This is an excerpt from an AVI movie generated by a targeting prototype tool courtesy of Mike Podlesak (DSTO) and Danny Gibbins and Steve Searle (CSSIP) for the ADF AIR 5418 programme.

The resulting 3x4 homogeneous matrix can be defined as a composite matrix given by a series of simpler transformations, from world (w) to platform (p) to gimbal (g) to camera (c) to image (i) or more succinctly: $M = M_{c2i} M_{g2c} M_{p2g} M_{w2p}$. These component matrices are defined as follows:

- $M_{c2i} = \begin{bmatrix} f_x & 0 & t_x \\ 0 & f_y & t_y \\ 0 & 0 & 1 \end{bmatrix}$

where $f_x = n_c / (2 \tan^{-1}(F_x / 2))$ and $f_y = n_r / (2 \tan^{-1}(F_y / 2))$ are the horizontal and vertical focal lengths given in pixels. Here the image centre, given by t_x, t_y is set to zero, however, they may be modified at a later point to incorporate registration

- $M_{g2c} = R_x(\frac{\pi}{2} - \phi_e) R_z \phi_a$ where $R_{xyz}(\alpha)$ are the rotations by the single angle α around the axes x, y, and z
- The M_{p2g} matrix represents the alignment of the gimbal with respect to the platform (the assumption here is that we have a pan/tilt camera). Obviously, this may change from flight to flight (and also, of course, during the flight) so it needs to be determined through a calibration technique¹⁵. Alternatively, as a default, it may be set to the identity matrix

- $M_{w2p} = [R | -Rt]$ where $R = R_y(-\theta_r) R_x(-\theta_p) R_z(\theta_h)$ and $t = \begin{bmatrix} x_p \\ y_p \\ z_p \end{bmatrix}$.

If we now have each pixel in the image as a direction vector or ray and if, a priori, the topology of the 3D scene is known (such as is provided by DTED¹⁶ or a digital terrain model) a ray from the image can be intersected with it to find the corresponding 3D point. If the scene surface is considered to be relatively flat this inverse mapping is simply a linear one.

Take p to be an image point and P to be its corresponding 3D scene point. We then have M (given above) as the projection matrix. If we assume that the surface of the scene can be approximated by a plane with normal n . We then have the following relationship:

$$P = \left[\left(I - \frac{mn^T}{m^T n} \right) M^{-1} \right] p$$

¹⁵ Calibrating cameras prior to tactical use allows for ease of calculating projective drift and other features of dynamic imaging such as terrain mapping. An initial set of calibration parameters also facilitates auto-calibration during flight which allows for the use of zooming, etc.

¹⁶ Digital Terrain Elevation Data. Developed by the US Department of Defense and the National Imagery and Mapping Agency. There are three levels of DTED data available depending on the location around the earth. DTED 0 is freely available on the internet. DTED 1-2 datasets are not generally available without specific permission from HQ NIMA. (<http://www.nima.mil>)

where M^- and m are, respectively, the pseudo-inverse and null vector of M , and I is the 4×4 identity matrix. If we assume a constant height above sea level, we have $n = (0, 0, -1, z_0)^T$. They also give the homogeneous inverse projection matrix:

$$\tilde{M} = \left(I - \frac{mn^T}{m^T n} \right) M^-$$

Registration is now used to improve the spatial accuracy of the technique. If we generate the approximate distance Z to the feature by using the raw projection matrix pan/tilt angle corrections can then be obtained by: $\theta_x = Zt_x / f_x$ and $\theta_y = Zt_y / f_y$. Complementary to this technique is the algorithm where the operator selects a feature for initialisation. This feature is then tracked during the video sequence allowing for finer calculation of M .

3.4 Video Registration

In a revisit of the body of work she so rigorously and usefully elucidated in 1992, Lisa Brown (now at IBM's T.J. Watson Research Centre in New York, USA) describes video registration as an addition to her earlier image registration review paper [14]. She describes three video registration categories, which we shall comment on, relative to our problem domain and taxonomy. These categories can be considered higher data dimensionality versions of the previously described Multimodal, Template, Viewpoint, and Temporal two-dimensional problems.

3.4.1 Video to Reference Imagery or 3D models: GeoLocation

In these applications, video streams are registered to reference imagery or 3D models (templates). This is an example of GeoLocation, or geo-registration, as described earlier in Section 3.3.2. Brown suggests that registration in this category is a form of template registration with high data dimensionality (where the templates are often 3D models and the input is the video stream). She also indicates that the higher dimensionality implies that often viewpoint and temporal registration issues are incorporated as well. A large number of algorithms available to solve these problems originate in the medical imaging domain where radiation treatment planning, fluoroscopy-guided procedures and computer-integrated surgery match 2D images with 3D models (which are derived from CT, MRI, SPECT, PET and other volumetric modalities which generate large 3D datasets).

3.4.2 Video to video registration

Brown describes this category as basically finding a video clip in a longer video sequence. A commercial example could be the detection of video copy, or pirating, detection. Two tactical examples relevant to our domain could be the *calibration* of multiple video cameras in UAVs or other low-flying, relatively slow aircraft. Also, the tactical alignment of a weaponised UAVs sensor to the three-dimensional surveillance survey space created from scout aircraft.

In the first example, a squadron of UAVs could align their sensors (video cameras) to the video feed generated from a high-flying command and control aircraft. Once calibrated, operations such as GeoLocation become more precise.

In the second example, a weapon on board a UAV could have its targeting mission simplified by the spatial alignment in a three-dimensional database created by numerous fly-overs by surveillance aircraft (or a *swarm* of smaller UAVs). Each video clip from the weapon's sensors, along with its metadata, could be matched to take advantage of the numerically dense survey *terrain cube* (as described previously). This could facilitate the identification of relocatable targets not recorded in previous imagery.

3.4.3 Frame to frame registration

This is the video equivalent of viewpoint registration as described earlier. The primary tactical application of problems in this area is the improvement of video quality. One example of this is the frame-to-frame registration needed to reduce the effect of engine vibration and sudden airframe movements as experienced by most small UAVs. Another critical example is the problem of video aliasing. In this case, the sensor may move between the first and second scans making GeoLocation and targeting in general, difficult.

4. Conclusions

This paper has included the formal definitions of image registration, described a proposed taxonomy, and described and elaborated two new algorithms: the *terrain cube* and the Hybrid registration method. These algorithms may be directly applied to relocatable targets and to the enhancement of mission planning capabilities in new weapons systems.

Special attention was paid to the process of abstraction as applied to the general registration problem. Video registration was added to the taxonomy to reflect the current and expected continued future needs for analysis of tactical (mostly UAV and missile-derived) imagery.

An algorithmic framework has also been described which can be applied to several current problems in tactical image registration. For example, GeoLocation is described to the level of implementation. This is an example of the different demands that tactical imaging requires from algorithms.

It should be noted that all the registration algorithms described here have assumed a rigid body. That is, there has been no allowance made for elastic or non-linear motion between images. Further work to extend the scope of this paper will include detailed assessment and categorisation of these elastic algorithms as well as the implementation of *terrain cubes* and Hybrid registration.

5. Acknowledgements

The author would like to thank Kevin Robinson for his careful reading of numerous versions of this paper and his helpful comments—especially those directing the ideas into military applications. Leszek Swierkowski and Barnaby Smith, the author's supervisor and manager respectively, provided encouragement and the time to write this paper and to gather ideas. This work has been done to support the research tasks in our group, which require accurate and relevant registration algorithms. These include air-to-ground targeting, superresolution, and other tactical applications.

6. References

- [1] Cuchet E., Knoploch J., Dormont D., Marsault C., *Registration in Neurosurgery and Neuroradiotherapy Applications*, **Proc. of Medical Robotics and Computer Assisted Surgery**, ed. Ayache N., pp. 31–38, Baltimore, MD, USA, November 1995.
- [2] Proceedings of **SPIE Vol. 4054** *Conference on Automated Geo-Spatial Image and Data Exploitation*.
- [3] Modersitzki, J., **Numerical Methods for Image Registration**, Oxford University Press, (2004).
- [4] Backman R.B., Spadaccini N., Wood F., *Differential Geometry Applied to Non-Rigid Registration*, Proc. of 1997 **The University of Western Australia Department of Computer Science Research Conference**, ed. Royle G.F., & Backman R.B., pp. 1–12, Yanchep, Western Australia, April 3–4, 1997.
- [5] Hartley, R.I. and Zisserman, A. **Multiple View Geometry in Computer Vision**, 2004", Cambridge University Press.
- [6] Ng, L, and Ibanez, L, **Insight into Images Principles and Practice for Segmentation, Registration, and Image Analysis**, A K Peters, (2004), ed. Yoo, T., pp. 239–306.
- [7] Backman, B., Vozzo, R., Swierkowski, L., *3D Terrain Information from UAV Imagery*, **Proceedings of Land Warfare Conference**, 27–30 September 2004, Melbourne, Australia, pp.337–344.
- [8] Brown, L. Gottesfeld, *A Survey of Image Registration Techniques*, **ACM Computing Surveys**, Vol. 24, No. 4, pp.325–376, December 1992.
- [9] Moons, T. and Frere, D. and Vandekerckhove, J. and Van Gool, L., *Automatic modeling and 3D reconstruction of urban house roofs from high resolution aerial imagery*, **Proc. European Conference on Computer Vision**, Vol 1, pp 410–425, 1998.

- [10] *Registration in Neurosurgery and Neuroradiotherapy Applications*, **Proc. of Medical Robotics and Computer Assisted Surgery**, ed. Ayache N., p. 33, Baltimore, MD, USA, November 1995.
- [11] Baillard, C. and Zisserman, A., *Automatic reconstruction of piecewise planar models from multiple views*, **Proc. IEEE Conference on Computer Vision and Pattern Recognition**, pp 559-565, 1999.
- [12] ARU87 Arun K.S., Huang T.S. and Blostein S.D., *Least-Squares Fitting of Two 3-D Point Sets*, **IEEE Transactions on Pattern Analysis and Machine Intelligence**, vol. 9, no. 5, pp. 698-703, September 1987.
- [13] Shekhar, C. and Chellappa, R., *Airborne Video Registration for Activity Monitoring in Video Registration*, Kluwer Academic Publishers, (2003), eds. Shah, M. and Kumar, R., pp.130-143.
- [14] Brown, L. {lisabr@us.ibm.com}, *Image Video Registration Retrospective*, in **Video Registration**, Kluwer Academic Publishers, (2003), eds. Shah, M. and Kumar, R., pp. 227-237.

DISTRIBUTION LIST

Registration for Tactical Imagery: An Updated Taxonomy

R. Bruce Backman

AUSTRALIA

DEFENCE ORGANISATION	No. of copies
Task Sponsor	
Chief WSD	1 Printed
S&T Program	
Chief Defence Scientist	1
Deputy Chief Defence Scientist Policy	1
AS Science Corporate Management	1
Director General Science Policy Development	1
Counsellor Defence Science, London	Doc Data Sheet
Counsellor Defence Science, Washington	Doc Data Sheet
Scientific Adviser Joint	1
Navy Scientific Adviser	Doc Data Sht & Dist List
Scientific Adviser – Army	1
Air Force Scientific Adviser	Doc Data Sht & Dist List
Scientific Adviser to the DMO	1
Systems Sciences Laboratory	
Research Leader	Doc Data Sht & Dist List
Head, Dr Barnaby Smith	1
Task Manager, Dr Leszek Swierkowski	1
Dr Farhan Faruqi	1
Author(s): Dr R Bruce Backman	1 Printed
Automation of the Battlespace Initiative	1
Head, Dr Anthony Finn	1
Trials Chief, Jean-Pierre Gibard	1
DSTO Library and Archives –	
Library Fishermans Bend	Doc Data Sheet
Library Edinburgh	2 printed
Defence Archives	1 printed

Capability Development Group

Director General Maritime Development	Doc Data Sheet
Director General Land Development – only for Army or Land reports	1
Director General Capability and Plans	Doc Data Sheet
Assistant Secretary Investment Analysis	Doc Data Sheet
Director Capability Plans and Programming	Doc Data Sheet

Chief Information Officer Group

Head Information Capability Management Division	Doc Data Sheet
AS Information Strategy and Futures	Doc Data Sheet
Director General Australian Defence Simulation Office	Doc Data Sheet
Director General Information Services	Doc Data Sheet

Strategy Group

Director General Military Strategy	Doc Data Sheet
Assistant Secretary Governance and Counter-Proliferation	Doc Data Sheet

Navy

Maritime Operational Analysis Centre, Building 89/90 Garden Island Sydney NSW	Doc Data Sht & Dist List
Deputy Director (Operations)	
Deputy Director (Analysis)	
Director General Navy Capability, Performance and Plans, Navy Headquarters	Doc Data Sheet
Director General Navy Strategic Policy and Futures, Navy Headquarters	Doc Data Sheet

Air Force

SO (Science) - Headquarters Air Combat Group, RAAF Base, Williamtown NSW 2314	Doc Data Sht & Exec Summary
---	-----------------------------

Army

SO (Science) - Land Headquarters (LHQ), Victoria Barracks NSW	Doc Data Sht & Exec Summary
SO (Science), Deployable Joint Force Headquarters (DJFHQ) (L), Enoggera QLD	Doc Data Sheet

Joint Operations Command

Director General Joint Operations	Doc Data Sheet
Chief of Staff Headquarters Joint Operations Command	Doc Data Sheet
Commandant ADF Warfare Centre	Doc Data Sheet
Director General Strategic Logistics	Doc Data Sheet

Intelligence and Security Group

Assistant Secretary Concepts, Capability and Resources	1
DGSTA Defence Intelligence Organisation	1
Manager, Information Centre, Defence Intelligence Organisation	1
Assistant Secretary Capability Provisioning, DSD	Doc Data Sheet
Director Advanced Capabilities, DIGO	Doc Data Sheet

Defence Materiel Organisation

Deputy CEO	Doc Data Sheet
Head Aerospace Systems Division	Doc Data Sheet
Head Maritime Systems Division	Doc Data Sheet
Program Manager Air Warfare Destroyer	Doc Data Sheet
CDR Joint Logistics Command	Doc Data Sheet
Guided Weapon & Explosive Ordnance Branch (GWEO)	Doc Data Sheet

OUTSIDE AUSTRALIA**INTERNATIONAL DEFENCE INFORMATION CENTRES**

US Defense Technical Information Center	1
UK Dstl Knowledge Services	1
Canada Defence Research Directorate R&D Knowledge & Information Management (DRDKIM)	1
NZ Defence Information Centre	1

SPARES	5 Printed
--------	-----------

Total number of copies: 31 Printed: 10 PDF: 21

DEFENCE SCIENCE AND TECHNOLOGY ORGANISATION DOCUMENT CONTROL DATA					
				1. PRIVACY MARKING/CAVEAT (OF DOCUMENT)	
2. TITLE Registration for Tactical Imagery: An Updated Taxonomy			3. SECURITY CLASSIFICATION (FOR UNCLASSIFIED REPORTS THAT ARE LIMITED RELEASE USE (L) NEXT TO DOCUMENT CLASSIFICATION) Document (U) Title (U) Abstract (U)		
4. AUTHOR(S) R. Bruce Backman			5. CORPORATE AUTHOR DSTO Defence Science and Technology Organisation PO Box 1500 Edinburgh South Australia 5111 Australia		
6a. DSTO NUMBER DSTO-TR-1855		6b. AR NUMBER AR-013-641		7. DOCUMENT DATE April 2006	
8. FILE NUMBER 2006/ 1069612		9. TASK NUMBER LRR 05/303		10. TASK SPONSOR CWSD	
				11. NO. OF PAGES 24	
				12. NO. OF REFERENCES 14	
13. URL on the World Wide Web http://www.dsto.defence.gov.au/corporate/reports/DSTO-TR-1855.pdf				14. RELEASE AUTHORITY Chief, Weapons Systems Division	
15. SECONDARY RELEASE STATEMENT OF THIS DOCUMENT <i>Approved for public release</i>					
OVERSEAS ENQUIRIES OUTSIDE STATED LIMITATIONS SHOULD BE REFERRED THROUGH DOCUMENT EXCHANGE, PO BOX 1500, EDINBURGH, SA 5111					
16. DELIBERATE ANNOUNCEMENT No Limitations					
17. CITATION IN OTHER DOCUMENTS Yes					
18. DSTO RESEARCH LIBRARY THESAURUS Algorithms, Image processing, Image registration, Taxonomies					
19. ABSTRACT Registration is central to Image Processing problems which use Tactical Imagery. Any application, which involves comparing two or more images, requires some type of Registration algorithm. These algorithms have evolved over the years and are generally grouped into three categories: 2D, 3D, and the reasonably unusual combination of 2D/3D. An updated classification (or taxonomy) for the diverse collection of algorithms is presented here and is described in detail. Also, two new algorithms are elucidated: the terrain cube and the Hybrid registration method. Many examples are given demonstrating the usefulness of this taxonomy and algorithms. The Medical Imaging field is the source for many of these examples, as numerous algorithms have their origin there. Complementary Military Imaging examples are also presented and described in terms of relevant platforms.					