Int J Software Informatics, Volume 6, Issue 3 (2012), pp. 381–398
International Journal of Software and Informatics, ISSN 1673-7288

E-mail: ijsi@iscas.ac.cn
http://www.ijsi.org
Tel: +86-10-62661040

Fusion of Data and Knowledge for Safe UAV Landing

Sally McClean, Bryan Scotney, Timothy Patterson, Philip Morrow, and Gerard Parr

(School of Computing and Information Engineering, University of Ulster,
Cromore Road, Coleraine, BT52 1SA, Northern Ireland)
Email: {si.mcclean,bw.scotney,pj.morrow,gp.parr}@ulster.ac.uk, patterson-t@email.ulster.ac.uk

Abstract Autonomous Unmanned Aerial Vehicles (UAVs) have the potential to significantly improve current working practices for a variety of applications including aerial surveillance and search-and-rescue. However before UAVs can be widely integrated into civilian airspace there are a number of technical challenges which must be overcome including provision of an autonomous method of landing which would be executed in the event of an emergency. A fundamental component of autonomous landing is safe landing zone detection of which terrain classification is a major constituent. Presented in this paper is an extension of the Multi-Modal Expectation Maximization algorithm which combines data in the form of multiple images of the same scene, with knowledge in the form of historic training data and Ordnance Survey map information to compute updated class parameters. These updated parameters are subsequently used to classify the terrain of an area based on the pixel data contained within the images. An image's contribution to the classification of an area is then apportioned according to its coverage of that area. Preliminary results are presented based on aerial imagery of the Antrim Plateau region in Northern Ireland which indicates potential in the approach used.
Key words: multi-resolution expectation maximization algorithm; UAV safe landing zone detection; UAV terrain classification

McClean S, Scotney B, Patterson T, Morrow P, Parr G. Fusion of data and knowledge for safe UAV landing. *Int J Software Informatics*, **Vol.6, No.3 (2012): 381–398. http://www.ijsi.org/1673-7288/6/i136.htm**

1 Introduction

Unmanned Aerial Vehicles (UAVs) provide many advantages over manned aircraft for a variety of military and civilian applications. Perhaps the most notable of these is the removal of humans from tasks which may be classified as dull, dangerous or dirty. For example, power line inspection[1], aerial surveillance[2] and monitoring of atmospheric pollution[3] respectively. Autonomous UAVs further enhance these advantages as they do not require real-time control by a human operator and therefore may have significantly lower operating costs.

This work is sponsored by The Engineering and Physical Sciences Research Council funded Sensing Unmanned Autonomous Aerial Vehicles project under Grants EP/F064217/1, EP/F064179/1, EP/F06358X/1 and through a Department for Employment and Learning studentship.
Corresponding author: Sally McClean, Email: si.mcclean@ulster.ac.uk
Received 2011-11-14; Revised 2012-04-04; Accepted 2012-09-04.

The Engineering and Physical Sciences Research Council (EPSRC) funded Sensing Unmanned Autonomous Aerial VEhicles (SUAAVE) project[4] is concerned with the creation of an autonomous system consisting of swarms of cooperating UAVs with an initial application scenario of mountain search-and-rescue. Currently the UAVs used within the SUAAVE project are Ascending Technologies Pelican platforms[5] (Fig. 1). These are a quadrotor design with a maximum flight speed of 50 km/h and a flight time of approximately 20 minutes. In addition to GPS and an inertial navigation system (INS) the Pelican is equipped with an ATOM processor board, an IEEE 802.11 networking card and a PointGray Chameleon colour camera. In many scenarios including mountain search-and-rescue swarms of autonomous, cooperating UAVs offer significant advantages over a single UAV operating in isolation. Not least is the potential to fuse observations from heterogeneous and geographically dispersed sensors to influence an efficient search strategy. A further advantage provided by utilising multiple UAVs is that in the case of a UAV failure mission execution can be continued by the remaining swarm members thereby achieving a level of robustness against node failure.

Figure 1. Ascending Technologies Pelican UAV platform

The United Kingdom (UK) Civil Aviation Authority's small UAV regulations[6] are currently similar to those specified for model aircraft. As such one stipulation is that a UAV must remain within 500 metres of the operator at all times. Intuitively this constraint greatly impacts upon the versatility and practical potential of the use of autonomous UAVs for mountain search-and-rescue. Before these regulations can be relaxed and UAVs fully integrated into unconstrained civilian airspace there are a number of technical challenges which must be addressed, including sense-and-avoid capabilities and provision of an autonomous safe landing system.

1.1 Background

As with manned aircraft UAVs are likely to encounter safety critical events such as prolonged loss of radio signal, for which the most prudent course of action may be to instruct the UAV to land as soon as possible (Fig. 2). From an operational safety standpoint it is not sufficient to command the UAV to land on the ground directly beneath it as this may present an unacceptable risk to animals, humans or property. Consequently one strand of research within the SUAAVE project has focused on autonomous Safe Landing Zone (SLZ) detection[7,8]. The approach consists of three main components, namely detection of potential SLZs, determining SLZ attribute

values and combining these values to assign an overall numeric safety score for each SLZ (Fig. 3). Perhaps the most critical stage of the autonomous SLZ detection process is accurately assigning the safety score to each potential landing site, as false positives and false negatives may have catastrophic consequences.

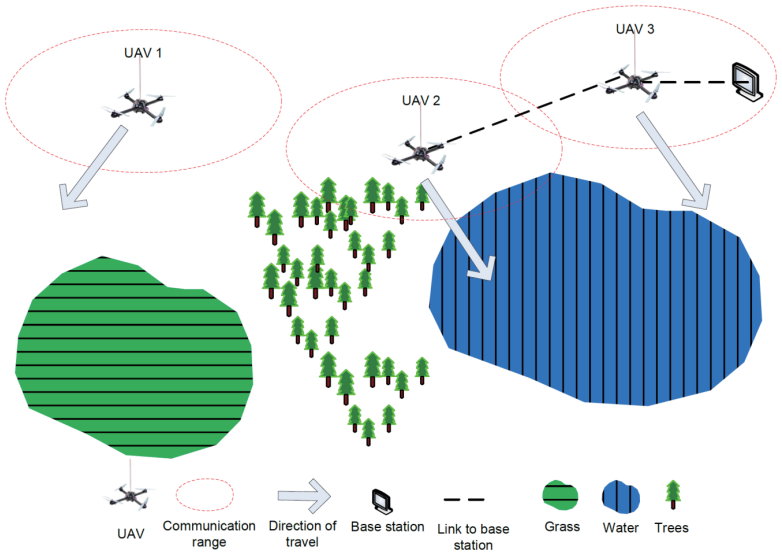

Figure 2. Example scenario where an emergency safe landing is required

The main objective of the SLZ detection algorithm is to minimize the risk of injuring persons, damaging property and where possible preserve the UAV and its payload. A secondary objective is to enable a fast and efficient retrieval of the UAV platform. With these objectives in mind a human expert assigns a suitability measure for landing to each type of terrain. This suitability measure in turn considerably influences the numeric safety score which a landing site receives. For example, in the majority of cases it can be assumed that grass is more suitable for landing on than water. Consequently it is vital to the success of the SLZ detection algorithm that the computed terrain classification of an area accurately reflects its real-world state. Within this paper we therefore focus upon the terrain classification component of SLZ detection (Fig. 3).

An example scenario where an emergency landing is required is depicted in Fig. 2. A key operational safety requirement is that all UAVs must remain within multi-hop communication range with the base station, thereby enabling commands such as *end mission* and *land immediately* to be transmitted. In this example a swarm of three UAVs are sent out to sense the environment in search of a missing person. Due to a GPS failure UAV1 navigates out of multi-hop communication range with the base station for a prolonged period of time. UAV1 subsequently implements the decision control process described in Ref. [7]. As the base station is not safely attainable without a reliable estimate of its real-world position UAV1 executes the SLZ detection algorithm described in Fig. 3 and determines that the terrain directly beneath it is a suitable landing site. UAV1 dynamically evaluates the chosen landing site as it descends and upon landing periodically transmits a beacon notifying other

UAVs of its presence should they fly within communication range.

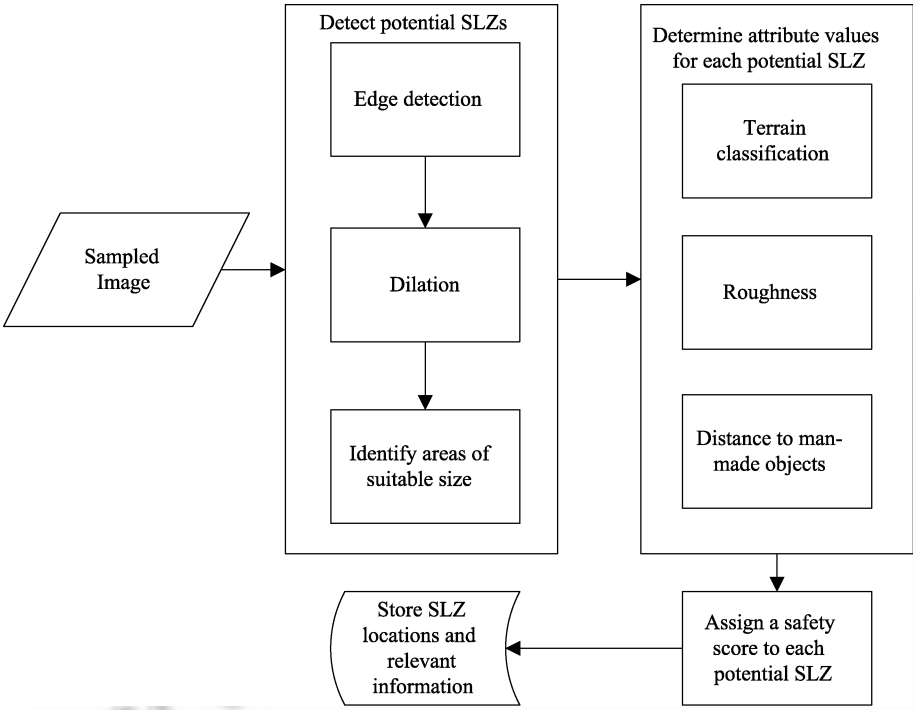

Figure 3. Autonomous SLZ detection algorithm overview

1.2 Motivation

The current approach to SLZ detection as discussed in Ref. [7] and Ref. [8] has focused on single images captured from a single UAV platform. Extending this work to incorporate multiple images captured by multiple UAVs has two main advantages for terrain classification:

1. The UAVs operate under strict power constraints and therefore are unlikely to observe an entire operational area. By combining observations such as terrain classification from neighbouring UAVs, a global interpretation of the operational area can be constructed which can enhance tasks such as SLZ detection.

2. The accuracy of computed terrain classification of an area is likely to be influenced by altitude of capture in addition to sensor type. Due to constraints imposed by payload capacity and available power it is likely that in an operational deployment each swarm member will only have a single sensing device. By enabling the incorporation of multiple images captured from multiple UAVs the potential is provided for observations from heterogeneous and geographically dispersed sensors to be used for terrain classification within SLZ detection.

An additional motivating factor behind this work is that, due to its inherently historic nature, training data for the supervised classification technique employed may not necessarily accurately reflect the spectral signatures of terrain classes within

the operational area. There are many reasons for this, including deviation in class spectral signatures between areas and variations due to seasonal changes.

Presented in this paper is a novel framework for the classification of terrain which combines knowledge in the form of class parameters estimated during an offline training phase and Ordnance Survey (OS) map information with data in the form of multiple images of an area. The images may have been captured at varying altitudes resulting in multiple resolutions, by different UAVs or at previous points in time. We extend our previous work on the Multi-Modal Expectation Maximization algorithm presented in Ref. [9] to ensure that both the knowledge and the data are combined and weighted in a principled fashion. This approach provides a real-time estimate of class parameters which are subsequently utilised to classify the terrain in the operational area.

The remainder of this paper is structured as follows; in Section 2 an overview of related work is presented. The algorithm and its various facets are discussed in Section 3. An evaluation based on preliminary results from high quality aerial images is presented in Section 4. Finally, conclusions and proposed further work are outlined in Section 5.

2 Related Work

"Information fusion is the study of efficient methods for automatically or semi automatically transforming information from different sources and different points in time into a representation that provides effective support for human or automated decision making."[10] There are numerous mathematical frameworks in the literature for data fusion tasks including fuzzy sets, possibility, rough sets, Dempster-Shafer and probability theory. For a comprehensive review of the current state-of-the-art the interested reader is referred to Ref. [11] and Ref. [12]. When considering data fusion within the context of terrain classification from aerial imagery two popular techniques include Dempster-Shafer and probability theory. The success of these techniques within this domain is in part due to their suitability for modelling uncertainty of membership within a well-defined class of objects[11].

2.1 Dempster-Shafer fusion

Dempster-Shafer theory of evidence provides a framework for combining sources of evidence to reach a degree of belief about the occurrence of an event in the presence of uncertainty and imprecision. Related applications which have utilised Dempster-Shafer theory include building detection[13] and terrain classification in satellite imagery[14].

In Ref. [13] work is presented in which Dempster-Shafer is used to fuse LIght Detection and Ranging (LIDAR) images with multi-spectral aerial images for classification of land cover with a particular emphasis on building detection. An influential component in Dempster-Shafer theory is the assignment of a mass for each class. In Ref. [13] a mass function is created which considers height differences between digital surface models created using LIDAR images and digital terrain models. Additionally, the mass function considers surface roughness, directedness of surface roughness and the normalized difference vegetative index computed from near infra-red and red portions in the light spectrum. The masses are subsequently created from the mass

function in a similar manner to fuzzy logic membership functions. Results are presented demonstrating a high accuracy of detection for large buildings; however, no discussion is provided about how the method performed when classifying the surrounding terrain.

One of the main advantages of using Dempster-Shafer theory is the ease with which ignorance, and therefore imprecision, in sensor readings can be quantified. For example, certain sensors may detect classes with greater reliability. A further advantage is that non-numerical data such as expert knowledge can be readily incorporated into the data fusion process. One such example is demonstrated in Ref. [14] where domain knowledge in the form of *if, then* rules is converted to masses and fused with probabilities from a Maximum Likelihood Classifier (MLC). An increase in classification accuracy of 6.8% is reported when using a MLC in conjunction with a knowledge base as opposed to a standalone MLC.

An early criticism of Dempster-Shafer theory was the computational complexity of reasoning. To accommodate this Guan and Bell[16] elaborated a method by Barnett[16] in which the problem space is partitioned and evidence subsequently clustered within these partitions. Guan and Bell consequently show that reasoning can be performed in linear time. However, a lingering concern is that Dempster's rule of combination may yield irrational results when confronted with conflicting data[17,18]. One could envisage such a scenario occurring within the application of mountain search-and-rescue where for example, two UAVs have conflicting beliefs about the terrain classification of an area.

2.2 Probabilistic fusion

Probabilistic data fusion, of which Bayes' theorem is a major constituent, has been widely used in real-world applications including fusing multiple sensor readings of the same scene for seabed classification[19] and detecting planar surfaces from UAV aerial imagery for autonomous SLZ detection[20].

The motivation behind the work in Ref. [19] is the hypothesis that classification accuracy will improve as the number of views increases. Multiple views of the same scene are fused using a joint likelihood calculation in conjunction with the assumption of independence between successive observations. Results are presented based on real-world, manually labelled seabed data which suggest increased classification accuracy for challenging terrain such as rippled seabed when viewed multiple times using differing sensor orientations. However, in the case of 'flat' seabed classification accuracy decreased from 78.5% when viewed once to 59.1% when viewed from four orientations, suggesting that multiple views do not always translate to increased classification accuracy.

One advantage of using Bayes' theorem for data fusion is that prior knowledge can be leveraged into the process in a principled fashion. Such prior knowledge exists for many applications including terrain classification from aerial imagery where map data may be available and incorporation of expert knowledge for medical diagnostics. Furthermore, as the outcome of the fusion process is a set of probabilities the amount by which further observations would improve confidence in the results can be quantified, for example by using the entropy of the results, as demonstrated in Ref. [19].

3 The Multi-Resolution Expectation Maximization Algorithm

The Multi-Resolution Expectation Maximization (MREM) algorithm is based upon work introduced in Ref. [9] where we proposed a model-based approach to image segmentation of the same scene acquired from multiple image modalities within a medical domain. Pixel intensity is modelled by a multivariate Gaussian distribution mixture in which components correspond to different data classes that represent image segments. The Expectation Maximization (EM) algorithm is adopted to estimate the maximum likelihood parameters that characterise the multivariate Gaussian distributions. The main advantage of our approach is that pixel data from all of the images are involved in estimating class probabilities, which in turn are combined with image specific data to give parameter estimations (means and variance-covariance matrices) for each class within an image.

In our current problem, images of a scene may have been captured at different times, varying altitudes or by heterogeneous sensor types. Our approach is based on well-established principles, namely probability models and maximum likelihood estimation. An underlying assumption is that the pixel intensity data within images are independent, whilst data from all images contribute to the probability of the relevant areas of a scene belonging to each class. We believe this assumption is a reasonable approximation and is in the spirit of, for example, Dempster's rule of combination, which also assumes that different sources of evidence are mutually independent. Additionally we assume a multivariate Gaussian distribution for p-dimensional pixel intensity data, where generally $p = 3$.

The search region is decomposed into a set of square grid cells A, of equal spatial resolution. The easting and northing position of the top left corner of each cell in A is known, resulting in a coordinate system which can be mapped to a real-world location and additionally provides each image with a common frame of reference. For each image it is assumed that an accurate easting and northing position of the top left corner for each pixel is available in addition to the spatial resolution.

3.1 Initialization

Our model-based segmentation uses model-based clustering to fit a mixture of multivariate Gaussian distributions that are characterised by parameters: mean vector $\boldsymbol{\mu}_k^{(I_r)}$ and variance-covariance matrix $\boldsymbol{\Sigma}_k^{(I_r)}$ for class $k = 1, ..., K$ and images I with resolutions $I_r = 1, ..., I_R$. Particularly in the multivariate case the initialization of these parameters can be influential in the quality of the solution provided by an EM based algorithm. There are two common strategies for initialization, namely random start positions and using the output from a k-means type algorithm[21].

Analogous to a k-means type approach we use an offline training phase to estimate initial values for the number of classes, K, and for each image resolution, respective $\boldsymbol{\mu}_k$ and $\boldsymbol{\Sigma}_k$ values. This training phase is conducted by a human expert familiar with the general operating area using, for example, satellite imagery. This approach to initialization does not necessarily constitute an additional overhead as knowledge in the form of training data is required to classify the output of the MREM algorithm for use within SLZ detection.

Further parameters are the prior class probabilities, denoted by π_k. In the absence of quantifiable knowledge about the prior probability of a class occurring equal

priors can be assumed. However, for the majority of regions within the UK knowledge about the operational area can be derived from data provided by the national mapping agency, Ordnance Survey (OS). Large scale, vector format OS data represents the UK landscape as a series of line, point and polygonal features. Each feature is specified in eastings/northings coordinates to an accuracy of ±0.4 metres[22]. Additional metadata is included such as the feature code field which can be considered as a form of terrain classification. Whilst this is historic in nature it can generally be assumed that the feature code field provides a reliable indicator of the terrain classification of an area.

For each top-level feature code a human expert is tasked with assigning prior class probabilities given the occurrence of that feature code. The prior probabilities incorporate human knowledge about the likelihood of a class given a vague or imprecise feature code, for example areas specified as vegetation may contain grass or forest. By utilising human knowledge in this way the likelihood of a change in classification may also be captured. There are many scenarios where it is envisaged that this could occur, for example during summer seasons rivers may dry up and a riverbed become an impromptu path, or new features may be constructed such as roads. An example assignment for a subset of classes and feature codes is outlined in Table 1.

Table 1 Example OS feature codes with assigned prior probabilities for a subset of classes

OS feature code	Assigned prior probabilities of class occurrence				
	Road	Path	Water	Grass	...
Road	0.6	0.2	0.04	0.04	...
Path	0.2	0.6	0.04	0.04	...
Water	0.04	0.1	0.4	0.1	...
Vegetation	0.05	0.1	0.05	0.3	...
...	...	...	...	...	...

3.2 Algorithm

We consider an image scene of an area in A for a given image consisting of S multidimensional pixels and containing K distinct classes labelled $C_1, ..., C_K$. The overall aim of the MREM algorithm is to assign each relevant cell in A to one of the K classes. The conditional probability density function (p.d.f.) of pixel intensity based on p-dimensional data $\mathbf{x}_s^{(I_r)}$ from pixel s of image resolution, I_r belonging to class k is then given by:

$$f_k^{(I_r)}(\mathbf{x}_s^{(I_r)}) = (2\pi)^{-p/2}\left|\mathbf{\Sigma}_k^{(I_r)}\right|^{-1/2} exp\{(\mathbf{x}_s^{(I_r)} - \boldsymbol{\mu}_k^{(I_r)})'\mathbf{\Sigma}_k^{(I_r)-1}(\mathbf{x}_s^{(I_r)} - \boldsymbol{\mu}_k^{(I_r)})\}. \quad (3.1)$$

If the corresponding class probabilities are denoted by $\pi_k, k = 1, ..., K$, then based on our assumption of independent images, we obtain the p.d.f. of pixel s as:

$$f^{(I_r)}(\mathbf{x}_s^{(I_r)}) = \sum_{k=1}^{K} \pi_k f_k^{(I_r)}(\mathbf{x}_s^{(I_r)}). \quad (3.2)$$

The MREM then proceeds as shown in Algorithm 1. Within each MREM iteration the E-Step is concerned with estimating posterior probabilities of each pixel in each image belonging to a class. In the M-Step classes and images each 'vote' to compute updated prior class probabilities. This enables both the classes and images to be weighted in an intuitive and principled manner as discussed in section 3.3. Additionally, during the M-Step updated $\boldsymbol{\mu}_k^{(I_r)}$ and $\boldsymbol{\Sigma}_k^{(I_r)}$ values are computed. Parameters estimated at the t^{th} iteration are denoted by a superscript (t).

Algorithm 1 MREM algorithm

1: Initialize the parameters using training data.
2: E-Step: compute the posterior probabilities:

$$P(C_k|\mathbf{x}_s^{(I_r)}) = P_{sk}^{(I_r)} = \frac{\pi_k^{(t)} f_k^{(I_r)}(\mathbf{x}_s^{(I_r)})}{\sum_{k=1}^{K} \pi_k^{(t)} f_k^{(I_r)}(\mathbf{x}_s^{(I_r)})}.$$

3: M-Step: update the parameters:

$$\pi_k^{(t+1)} = \frac{1}{\sum_{I_r=1}^{I_R} I_{rS}} \sum_{I_r=1}^{I_R} \sum_{s=1}^{S} P_{sk}^{(I_r)},$$

$$\boldsymbol{\mu}_{I_{rk}}^{(t+1)} = \frac{\sum_{s=1}^{S} \mathbf{x}_{I_{rs}} P(C_k|\mathbf{x}_{I_{rs}})}{\sum_{s=1}^{S} P(C_k|\mathbf{x}_{I_{rs}})},$$

$$\boldsymbol{\Sigma}_{I_{rk}}^{(t+1)} = \frac{\sum_{s=1}^{S} P(C_k|\mathbf{x}_{I_{rs}})((\mathbf{x}_{I_{rs}} - \boldsymbol{\mu}_{I_{rk}}^{(t+1)})^T(\mathbf{x}_{I_{rs}} - \boldsymbol{\mu}_{I_{rk}}^{(t+1)}))}{\sum_{s=1}^{S} P(C_k|\mathbf{x}_{I_{rs}})},$$

Update $f_k^{(I_r)}(\mathbf{x}_s^{(I_r)})$ using parameters estimated at iteration $^{t+1}$.

4: Compute termination criterion for example, (pseudo) likelihood.
5: Repeat steps 2 to 4 iteratively until convergence is attained.
6: Segmentation step: assign the class label for relevant cells in A:

$$A_{e,n} = \arg\max_k \left(\frac{1}{I_R} \sum_{I_r=1}^{I_R} P_{sk}^{(I_r)}\right).$$

The primary output of the MREM algorithm is a terrain classification for each observed cell in A. A secondary output is an updated estimate of class parameters. Whilst it may appear appropriate to utilise the updated class parameters to initialize the MREM algorithm upon each execution this may not necessarily increase the accuracy of the approach. For example, brightly coloured objects such as flowers may skew the estimated parameters of the nearest class. Therefore, in this work class parameters are initialized using training data upon each execution of the MREM algorithm. This ensures that whilst a measure of flexibility is incorporated to reflect the historic nature of the training data the global impact of updated class parameters is limited. Online learning of class parameters in an outdoor environment will form a beneficial, albeit challenging part of future work.

3.3 Weighting of images

Within the context of terrain classification to assist in the assignment of a SLZ safety score[8] it is likely that images will vary in terms of their reliability and precision. It is therefore desirable to weight an image's 'vote'. Three attributes which would influence the weighting an image receives are the altitude of capture, time since capture and sensor type. The general concept and how the weights would be incorporated into the MREM algorithm are outlined below. In an actual implementation the influence that the specific attribute values would have upon an image's weighting could be determined empirically.

In the general case the altitude of image capture can be influential upon the terrain classification assigned to an area. Images captured at a high altitude will be of a lower spatial resolution and therefore contain less information about a real world area than their lower altitude counterparts. However, it cannot be assumed that a lower altitude of capture will necessarily translate to a higher quality result. For example at a very low altitude grass may exhibit similar spectral signatures to forest. Furthermore valuable contextual information may not be available when images are captured at a very low altitude. It is therefore desirable to weight the altitude of capture between a lower and upper bound.

Many of the terrain classes likely to be encountered during a mountain search-and-rescue have the potential to change throughout the duration of the mission. There are two main types of possible changes. Firstly a class may be to some extent characteristically dynamic. One such example is water, which may be influenced by precipitation and evaporation. Secondly, a class may have a higher likelihood of containing external objects, i.e., objects which are not considered as terrain. For example, the terrain class road which incorporates spectrally and functionally similar areas such as car parks may change frequently resulting in an area which previously may have been considered as a potential SLZ becoming wholly unsuitable for landing in. Similarly the contents of rural areas such as fields may change due to animals grazing. It is therefore desirable to decrease the weight which an image has upon the classification of an area as its time since capture increases.

Whilst the primary sensing device utilised within the SUAAVE project is a COTS colour camera it is advantageous to include a mechanism whereby an image is weighted according to the type and quality of the capturing sensor. This enables knowledge such as a sensor's aptitude for detecting certain classes to be incorporated. For example, in Ref. [23] it is shown that thermal infra-red cameras can be very successful at detecting wet areas as they are generally cooler than the surrounding vegetation.

The formulation of step 3 in the MREM algorithm allows the efficient and lightweight incorporation of weights into the approach. Let w_{1Irk} be the weighting associated with altitude of capture for class k of image resolution I_r. Similarly let w_{2Irk} and w_{3Irk} be the weightings associated with time since capture and sensor type respectively. The combined weighting for each class in each image is then calculated as,

$$w_{Irk} = \frac{w_{1Irk} + w_{2Irk} + w_{3Irk}}{\sum_{k=1}^{K} \sum_{j=1}^{3} w_{jmk}}. \tag{3.3}$$

These weights are incorporated by modifying the relevant part of step 3 in the MREM

algorithm to,

$$\pi_k^{(t)} = \frac{1}{\sum_{I_r=1}^{I_R} I_{rS}} \sum_{I_r=1}^{I_R} \sum_{s=1}^{S} w_{Irk} P_{sk}^{(I_r)}. \tag{3.4}$$

In a real-world implementation it is envisaged that the actual weightings associated with attribute values for altitude of capture, time since capture and sensor type would be learned from a human expert during an offline training phase. Whilst this will form part of future work, possible approaches include weighting an image inversely by time since capture and weighting image credibility based on intraclass variance of spectral values. When used in conjunction with a threshold such a weighting mechanism may be further used to assist in excluding potentially unuseful images from consideration by the MREM algorithm.

3.4 Incorporating images of varying spatial resolution

For the application scenario of combining aerial images captured by a UAV with additional images during a search-and-rescue mission it is unrealistic to assume that each image has the same pixel scale. Therefore an image's contribution towards the terrain classification of an area is apportioned according to its coverage of that area.

When considering images of varying spatial resolution with respect to the spatial resolution of A there are three possibilities, namely, higher resolution, equal resolution or lower resolution. Similar steps are taken to compute the coverage for each possibility. As an example, the case where an image has a higher resolution is discussed in the following subsection.

3.4.1 Incorporating images of higher spatial resolution

Where any image has a higher spatial resolution than A, i.e. a single pixel in I covers a smaller real-world area than a cell in A, there are nine different types of coverage as outlined in Fig. 4. In order to calculate the coverage of a pixel in I of a cell in A the type of coverage must first be detected. The coverage in the x axis, c_x and y axis, c_y can then be subsequently calculated. To compute the coverage the easting, A_e, northing, A_n coordinates of the top left corner of the relevant cell in A in addition to the easting, I_e, northing, I_n coordinates of the top left corner of the pixel under consideration is necessary. The spatial resolution of A denoted by A_r and I denoted by I_r is also required. The steps in determining the type and area of coverage of a single pixel in I are delineated below:

(i) Case 1 occurs when the x-axis of a pixel at $I_{e,n}$ is partly to the left of $A_{e,n}$ and the y-axis is partly above (Fig. 4i). This case is detected when:

$$I_n > A_n \text{ and } I_n - I_r < A_n \text{ and } I_e < A_e \text{ and } I_e + I_r > A_e,$$

c_x, c_y can then be calculated as,

$$c_x = I_e - A_e + I_r, c_y = A_n - (I_n - I_r). \tag{3.5}$$

(ii) Case 2 occurs when the x-axis of a pixel at $I_{e,n}$ is fully contained within $A_{e,n}$ and the y-axis is partly above (Fig. 4ii). This case is detected when:

$$I_n > A_n \text{ and } I_n - I_r < A_n \text{ and } I_e \geqslant A_e \text{ and } I_e + I_r \leqslant A_e + A_r,$$

c_x, c_y can then be calculated as,

$$c_x = I_r, c_y = A_n - (I_n - I_r). \tag{3.6}$$

(iii) Case 3 occurs when the x-axis of a pixel at $I_{e,n}$ is partly to the right of $A_{e,n}$ and the y-axis is partly above (Fig. 4iii). This case is detected when:

$$I_n > A_n \text{ and } I_n - I_r < A_n \text{ and } I_e > A_e \text{ and } I_e + I_r > A_e + A_r,$$

c_x, c_y can then be calculated as,

$$c_x = A_e + A_r - I_e, c_y = A_n - (I_n - I_r). \tag{3.7}$$

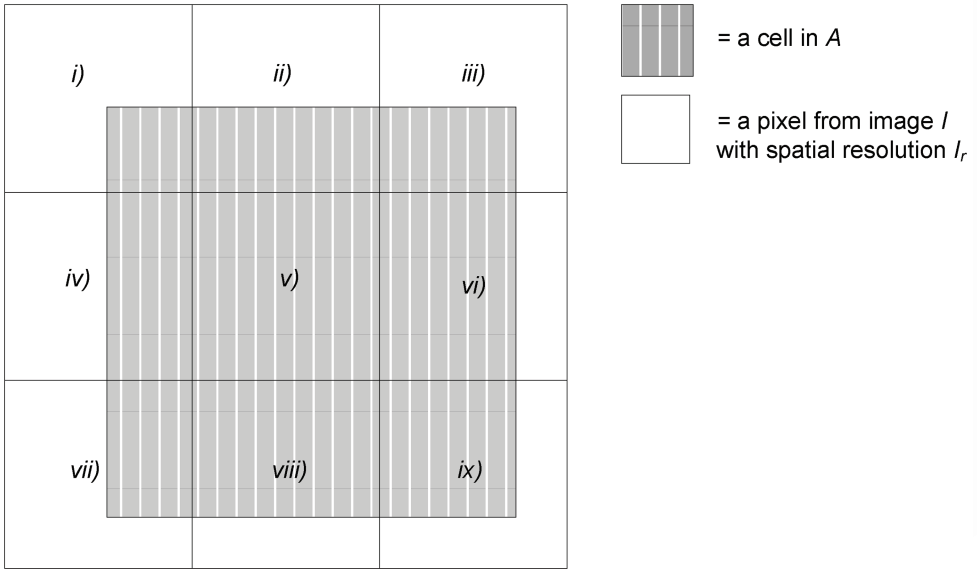

Figure 4. Possible types of coverage where an image has a higher spatial resolution than A

(iv) Case 4 occurs when the x-axis of a pixel at $I_{e,n}$ is partly to the left of $A_{e,n}$ and the y-axis is fully contained within it (Fig. 4iv). This case is detected when:

$$I_n \leqslant A_n \text{ and } I_n - I_r \geqslant A_n - A_r \text{ and } I_e < A_e \text{ and } I_e + I_r > A_e,$$

c_x, c_y can then be calculated as,

$$c_x = I_e - A_e + I_r, c_y = I_r. \tag{3.8}$$

(v) Case 5 occurs when the entire pixel at $I_{e,n}$ is fully contained within a cell at $A_{e,n}$ (Fig. 4v). This case is detected when:

$$I_n \leqslant A_n \text{ and } I_n - I_r \geqslant A_n - A_r \text{ and } I_e \geqslant A_e \text{ and } I_e + I_r \leqslant A_e + A_r,$$

c_x, c_y can then be calculated as,

$$c_x = I_r, c_y = I_r. \tag{3.9}$$

(vi) Case 6 occurs when the x-axis of a pixel at $I_{e,n}$ is partly to the right of $A_{e,n}$ and the y-axis is fully contained within it (Fig. 4vi). This case is detected when:

$$I_n \leqslant A_n \text{ and } I_n - I_r \geqslant A_n - A_r \text{ and } I_e > A_e \text{ and } I_e + I_r > A_e + A_r,$$

c_x, c_y can then be calculated as,

$$c_x = A_e + A_r - I_e, c_y = I_r. \tag{3.10}$$

(vii) Case 7 occurs when the x-axis of a pixel at $I_{e,n}$ is partly to the left of $A_{e,n}$ and the y-axis is partly below (Fig. 4vii). This case is detected when:

$$I_n < A_n \text{ and } I_n - I_r < A_n - A_r \text{ and } I_e < A_e \text{ and } I_e + I_r > A_e,$$

c_x, c_y can then be calculated as,

$$c_x = I_e - A_e + I_r, c_y = I_n - (A_n - A_r). \tag{3.11}$$

(viii) Case 8 occurs when the x-axis of a pixel at $I_{e,n}$ is fully contained in $A_{e,n}$ and the y-axis is partly below (Fig. 4viii). This case is detected when:

$$I_n < A_n \text{ and } I_n - I_r < A_n - A_r \text{ and } I_e \geqslant A_e \text{ and } I_e + I_r \leqslant A_e + A_r,$$

c_x, c_y can then be calculated as,

$$c_x = I_r, c_y = I_n - (A_n - A_r). \tag{3.12}$$

(ix) Case 9 occurs when the x-axis of a pixel at $I_{e,n}$ is partly to the right of $A_{e,n}$ and the y-axis is partly below (Fig. 4ix). This case is detected when:

$$I_n < A_n \text{ and } I_n - I_r < A_n - A_r \text{ and } I_e > A_e \text{ and } I_e + I_r > A_e + A_r,$$

c_x, c_y can then be calculated as,

$$c_x = A_e + A_r - I_e, c_y = I_n - (A_n - A_r). \tag{3.13}$$

Having calculated the coverage of a pixel at $I_{e,n}$ of a cell at $A_{e,n}$ the influence, ι, on the terrain classification at $A_{e,n}$ can be apportioned according to the coverage by,

$$\iota = \frac{c_x \times c_y}{A_r^2}. \tag{3.14}$$

Thus step 6 of the MREM algorithm becomes,

$$A_{e,n} = \arg\max_k \left(\frac{1}{I_r} \sum_{I_r=1}^{I_R} \iota P_{sk}^{(I_r)}\right), \tag{3.15}$$

resulting in pixels which cover a larger area of a cell in A having a greater influence upon the final classification.

3.5 Multiple UAVs

The algorithm described in this paper provides a method whereby multiple images of the same scene each contribute to the classification of an area. Whilst the steps executed are the same regardless of the image source in theory these multiple images may have been captured at different points in time, at varying altitudes or by different members of a UAV swarm (Fig. 2). However in practice there are a number of issues which must be considered before the incorporation of images from different members of the UAV swarm can be implemented in a real-world scenario.

Within a swarm based sensing platform such as multiple UAVs the question of when to transmit and receive information such as images can be of fundamental importance to the overall success of the mission. It is infeasible to assume unlimited bandwidth especially bearing in mind that other tasks requiring communication such as path planning and collision avoidance are simultaneously executing. Thus the consideration of when to share sensed images will be based upon the number of hops required, available bandwidth and the time required for transmission, in addition to the relevant UAV system attribute states, for example battery life.

An influential parameter when considering time for transmission is the volume of data to be transmitted. As the MREM algorithm utilises the actual spectral values contained within the images to estimate new class parameters it is a requirement that the complete image is transmitted along with an associated weighting. However in a real-world implementation these images would be compressed, thus reducing the bandwidth requirement. As the UAVs have finite battery life one possible approach to limiting the required bandwidth is to prioritise which images are transmitted based on areas that a UAV has the potential to navigate to. In the event of a UAV requesting images from surrounding swarm members for assistance in locating a safe landing zone a further method of prioritization may be to transmit only images of areas which have suitable terrain for landing in. Determining the circumstances under which image data are transmitted will be the subject of further research.

4 Evaluation

A preliminary evaluation of the MREM algorithm presented in this paper was conducted using commercially available high quality aerial imagery of the Antrim Plateau region in Northern Ireland. Whilst this aerial imagery was captured during manned flights it is spectrally representative of images sensed by a UAV. One of the main advantages of using such aerial imagery at this stage of evaluation is that it generally does not contain noise such as clouds. A further advantage is that it is geo-registered thus enabling OS data to be readily incorporated. As a case study aerial imagery of two areas, each measuring approximately 0.5km by 0.25km were chosen. These areas contained 9 main classes (road, path, water, shadow, trees, soil, grass, scrubland and gorse) which are similar to those likely to be encountered by a UAV performing a search-and-rescue mission in mountainous terrain. At this stage in development a distinct class is used for 'shadow' as due to its spectral appearance accurately determining the underlying terrain type is challenging, however may form part of future work.

For the purposes of this evaluation three spatial resolutions were used to simulate varying UAV altitudes. One low resolution (1 pixel = 1 metre) image was used to simulate an image captured at a high altitude; four images of medium spatial

resolution (1 pixel = 0.25 metre) were used to simulate images captured at a medium altitude and sixteen images of high spatial resolution (1 pixel ≈ 0.1 metre) used to simulate images captured at a low altitude.

4.1 Accuracy

To determine the accuracy of the MREM algorithm portions of the study area were manually labelled by a human expert using a Matlab based GUI. Three experiments were conducted which demonstrate how additional image resolutions influence the overall classification. The first experiment was based on a single low resolution image. The second experiment combined both the low resolution image with the four medium resolution images. Finally a third experiment combined all the image resolutions. The results of these experiments can be found in Fig. 5.

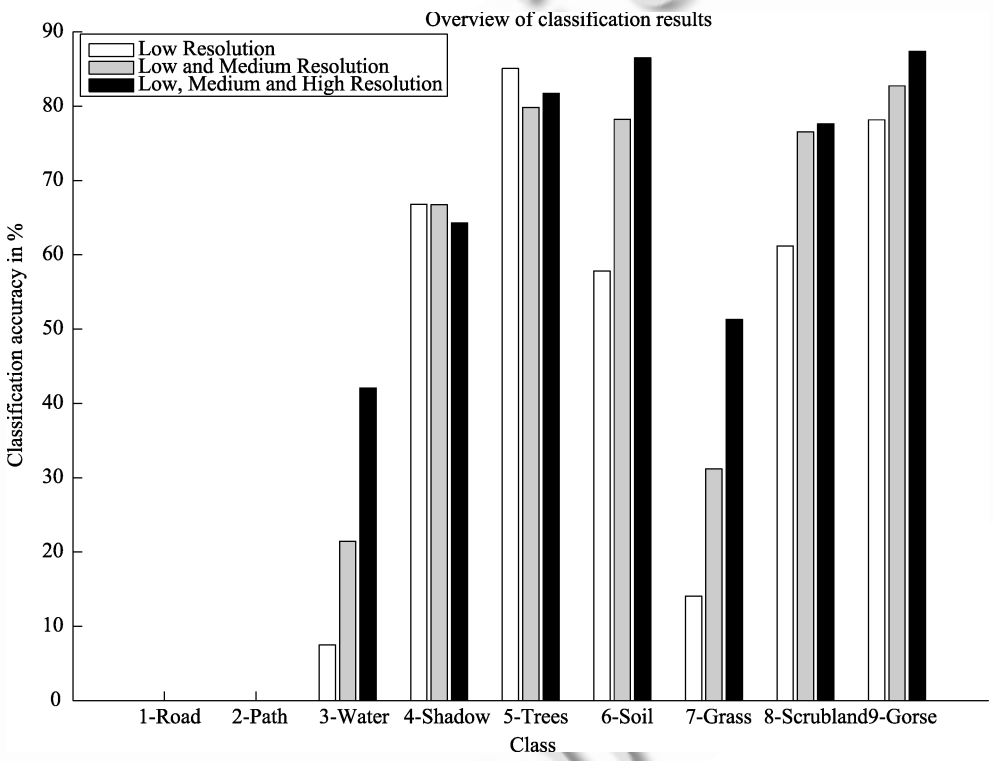

Figure 5. Classification results demonstrating the effect of additional resolutions upon accuracy

There are two main reasons for the low classification accuracy of road and path. Firstly, these classes account for a very small portion of the study area and are therefore assigned a low likelihood of occurring firstly by initial weights based on OS data and subsequently by the MREM algorithm. Secondly these classes have low spectral separability with other classes when considered in the RGB colour space, for example a path is visually similar to scrubland and road is visually similar to shadow and water. In many cases an indicator of man-made features such as roads or paths can be the presence of geometrical features such as parallel lines. These lines can be recognized in aerial imagery using image processing techniques such as edge detection

and edge linking. It is expected that incorporating such knowledge into the MREM algorithm would increase the classification accuracy of these classes. The approach to weighting discussed in this paper naturally lends itself towards giving a single class within an image a higher weight. For example, roads, paths and water may be more accurately detected by higher resolution imagery than grass. Determining these weights in a principled fashion using real data from a UAV will form an invaluable part of future work.

When considering both study areas the overall classification accuracy increased from 40.38% when classified using low resolution imagery to 56.59% when classified using both low and medium resolution imagery. When the areas were classified using three resolutions the overall classification accuracy further increased to 67.58% suggesting that additional image resolutions can lead to increased accuracy. Whilst for the application of terrain classification to assist in UAV SLZ detection these overall accuracies are relatively low possible approaches which may increase the classification accuracy may be to consider alternative colour spaces and to utilise a textural based classifier as opposed to the 'per-pixel' approach used in this work. In particular it is envisaged that textural based classification techniques would enhance the classification accuracy of grass as a pixel's neighbourhood is considered. Additionally, post classification processing, for example using Markov Random Fields may further increase the accuracy.

4.2 Computational cost

The MREM algorithm was implemented in Matlab and evaluated on an Intel Core 2 Duo 3Ghz desktop PC with 3GB RAM. The termination criterion for the MREM algorithm was that of convergence in estimated class parameters or a change in the log likelihood of less than 0.001. A summary of the required time is outlined in Table 2. On average the MREM algorithm required 4 iterations before termination. When considering three image resolutions this translates to an average time of 27.2 seconds which represents a significant time overhead when considered within the context of SLZ detection for a UAV. It is however hoped that with further refinement and optimization of code this may be reduced to an average time of 1 second per image resulting in 4 iterations of three images, taking approximately 12 seconds.

Table 2 Average values for the MREM algorithm per iteration

Experiment	# of images considered	Time(s)
Low res	1	0.6
Low and medium res	2	3.6
Low, medium and high res	3	6.8

5 Conclusions and Further Work

Presented in this paper is a novel framework which combines knowledge in the form of class parameters estimated during an offline training stage and OS map information with data in the form of multiple images of a scene to classify the terrain of an area. The MREM algorithm is used to update class parameters based on the spectral values contained within each image resolution. Updated class parameters are then used to classify each image and its influence upon the classification of an area apportioned according to its coverage of that area. A number of areas of future work have been identified which will further the usefulness of this framework for online terrain classification applications such as SLZ detection for UAVs. Initial results are presented based on spectrally representative aerial images of the Antrim Plateau Region in Northern Ireland which indicate potential in the approach used.

Perhaps one of the main advantages of the approach discussed in this paper is the ability to efficiently incorporate weights into the terrain classification process. These weights may be based on various attributes including altitude of capture, sensor type or time since capture. For a real-world safety critical application it is not sufficient to assign these weights arbitrarily and therefore an important part of future work will be analysing how values for each of these attributes impact upon the reliability of the terrain classification provided by an image.

Currently a per-pixel based approach is used for the terrain classification component of the MREM algorithm. For classes with high spectral separability this approach provides a relatively high rate of classification accuracy. However, it can yield poor results for classes which are visually similar. An important part of future work will therefore be to implement textural based classification methods in addition to a post classification component using, for example Markov Random Fields.

A preliminary evaluation of the framework presented in this paper was performed using high quality aerial imagery. An immediate extension of this work is therefore to implement the MREM algorithm on a UAV platform. This will introduce a number of constraints such as the limited computational power, memory capacity and available battery life however will enable us to further validate the approach. It is expected that with the implementation of the further work outlined in this section the MREM algorithm will prove to be a valuable component within the SLZ detection process.

Acknowledgements

This research was supported by a Department for Employment and Learning studentship and through the Sensing Unmanned Autonomous Aerial Vehicles (SUAAVE) project under grants EP/F064217/1, EP/F064179/1 and EP/F06358X/1.

References

[1] Li Z, Liu Y, Walker R, Hayward R, Zhang JN. Towards automatic power line detection for a UAV surveillance system using pulse coupled neural filter and an improved Hough transform. Machine Vision and Applications, September 2009, 21(6): 677–686.

[2] Jakob M, Semsch E, Pavlicek D, Pechoucek M. Occlusion-aware multi uav surveillance of multiple urban areas. Proc. of the 9th International Conference on Autonomous Agents and Multiagent Systems. Foundation for Autonomous Agents and Multiagent Systems. Toronto, Canada. 2010. 1407–1408.

[3] Ramana V, Ramanathan M, Nguyen V, Xu H, Pistone K, Corrigan C, Feng Y, Zhue A, Kim S, Yoon S, Carmichael G, Schauer J. Using unmanned aircraft to measure the impact of pollution plumes on atmospheric heating rates and cloud properties. American Geophysical Union. Fall Meeting. 2009.

[4] Cameron S, Parr G, Nardi R, Hailes S, Symington A, Julier S, Teacy L, McClean S, McPhillips G, Waharte S, Trigoni N, Ahmed M. SUAAVE: Combining Aerial Robots and Wireless Networking. Unmanned Air Vehicle Systems, number 01865, Bristol. 2010. 7–20.

[5] Ascending Technologies. AscTec Falcon, Accessed September 2011.

[6] Haddon DR, Whittaker CJ. UK-CAA policy for light UAV systems. UK Civil Aviation Authority, London, 2004.

[7] Patterson T, McClean S, Parr G, Morrow P, Teacy L, Nie J. Integration of Terrain Image Sensing with UAV Safety Management Protocols (In Press). The 2nd International ICST conference on Sensor Systems and Software, Florida, USA, 2010. Springer.

[8] Patterson T, McClean S, Morrow P, Parr G. Towards autonomous safe landing site identification from colour aerial images (In Press). 2010 Irish Machine Vision and Image Processing Conference, Limerick, Ireland, IEEE. 2010.

[9] Hong X, McClean S, Scotney B, Morrow P. Model-Based Segmentation of Multimodal Images. Springer Berlin, 2007. 604–611.

[10] Boström H, Andler SF, Brohede M, Johansson R, Karlsson A, Van Laere J, Niklasson L, Nilsson M, Persson A, Ziemke T. On the defination of information fusion as a field of research. School of Humanities and Informatics, University of Skövde, Tech. Rep., 2007.

[11] Khaleghi B, Khamis A, Karray FO. Multisensor data fusion: a review of the state-of-the-art (In Press). Information Fusion, August 2011.

[12] Dubois D, Liu W, Ma J, Prade H. A principled discussion of information combination rules in different representation settings. Soft Computing and Pattern Recognition (SoCPaR). 2011 International Conference of, Dalian. IEEE. 2011. 446–451.

[13] Rottensteiner F, Trinder J, Clode S, Kubik K. Using the Dempster-hafer method for the fusion of LIDAR data and multi-spectral images for building detection. Information Fusion. December 2005, 6(4): 283–300.

[14] Al Momani B, McClean S, Morrow P. Using DempsterShafer to incorporate knowledge into satellite image classification. Artificial Intelligence Review. 2007, 25(1-2): 161–178.

[15] Guan JW, Bell DA. Evidence Theory and its Applications: Volume 1. Elsevier, North-Holland, 1991.

[16] Barnett JA. Computational methods for a mathematical theory of evidence. Seventh Int. Joint Conf. Artificial Intelligence, Vancouver, B.C., 1981. 868–875.

[17] Zadeh LA. Review of "a mathematics theory of evidence". AI Magazine, 1984, 5(3): 81–83.

[18] Liu W. Analyzing the degree of conflict among belief functions. Artificial Intelligence, August 2006, 170(11): 909–924.

[19] Williams DP. Bayesian data fusion of multiview synthetic aperture sonar imagery for seabed classification. IEEE transactions on image processing : a publication of the IEEE Signal Processing Society, June 2009, 18(6): 1239–54.

[20] Bosch S, Lacroix S, Caballero F. Autonomous detection of safe landing areas for an UAV from monocular images. IEEE/RSJ International. Beijing, China, IEEE. 2006. 12–18.

[21] Biernacki C, Celeux G. Choosing starting values for the EM algorithm for getting the highest likelihood in multivariate Gaussian mixture models. Computational Statistics & Data Analysis. 2003, 41: 561–575.

[22] Ordnance Survey Northern Ireland. OSNI Large-Scale Technical Specification, Accessed September 2011.

[23] Rankin AL, Matthies LH. Passive sensor evaluation for unmanned ground vehicle. Journal of Field Robotics, 2010, 27(4): 473–490.