Behaviour Recognition for Spatially Unconstrained Unmanned Vehicles

Rolf Baxter, David Lane, Yvan Petillot
Ocean Systems Laboratory, Heriot-Watt University
Edinburgh, Scotland, EH14 4AS
{R.Baxter, D.M.Lane, Y.R.Petillot} @hw.ac.uk

Abstract

We use trajectory based techniques to perform location independent behaviour recognition on an unmanned underwater vehicle. Unmanned vehicles have applications in both surveillance and structural inspection, but require robust, location independent behaviour recognition. Previous research has used GPS or location based states that tie model parameters to specific locations, and have rarely considered performance under specific levels of noise. We use location independent action based states to recognise high level behaviour using flat Hidden Markov Models. We validate this approach by comparing performance under different levels of noise, achieving 78% classification precision under 50% corruption.

1 Introduction

Inferring the intentions of an agent is often posed as Behaviour Recognition, in which the observable actions of an agent are used to infer its internal state. In Activity Recognition the objective is to recognise low-level behaviours such as making a snack, while Plan Recognition extends this by recognising the hierarchical structure linking behaviours.

Previous research in behaviour recognition has included visual surveillance [Nguyen et al., 2005] and GPS tracking [Liao et al., 2007a], but is often constrained to recognising repeatable trajectories within known environments. This paper motivates applications in the new domains of Unmanned Vehicles (UxVs) in which there are several additional challenges. The increasing use of these vehicles has spawned a new interest in predicting the intentions of the agents they observe, and of the vehicles themselves.

In military domains unmanned vehicles are increasingly used to provide access to difficult areas. Aerial applications include covert visual surveillance, in which up to 40 hours of video[1] are collected per flight. [Oliver et al., 2002] highlight behaviour recognition as a critical step in automating visual surveillance, in which its application would automate processing large volumes of captured data. This complements human analysis by identifying threats within busy and noisy environments. A further benefit is the potential for increased autonomy, allowing vehicles to autonomously detect activity of interest for more detailed inspection.

Other applications of unmanned vehicles provide access to the oceans. Early vehicles were operated remotely using communications tethers, although more recently, their replacement with autonomy has allowed both shallow and deep water missions. Applications in the civil domain include ocean floor survey and pipe inspection, while military applications include sea-mine countermeasures.

As the autonomy of unmanned vehicles increases additional benefits may be gained through on-board behaviour recognition. For example, recent research in UxV collaboration has allowed common sub-goals to be shared between multiple vehicles [Sotzing et al., 2007]. Using a limited prediction model, collaboration is able to continue under the loss of communications, but is limited by outage duration and knowledge of shared sub-goals. On-board behaviour recognition would complement this work by increasing reliability under the prolonged loss of inter-vehicle communications. This is especially beneficial for re-planning systems, in which the current sub-goals of one agent may be unknown to another. Behaviour recognition would allow these systems to continue safe collaboration through behaviour observation.

In adversarial domains the actions of others can interfere with goal completion. One solution to this is to predict the intentions of opponents, allowing on-line re-planning to respond to perceived threats. As an example, an autonomous ground vehicle performing perimeter surveillance might identify unknown persons within its vicinity. In a military domain, a person walking along a road may require the vehicle to remain still, while a hostile formation approaching the vehicle may require transmitting a location signal. In either scenario, re-planning can only take place upon identifying the adversary's intentions.

The main research contribution in this paper is evaluating the feasibility of a Hidden Markov Model approach for UxV sub-goal recognition. We continue to use the term sub-goal recognition to highlight our ultimate goal of plan recognition. Previous applications of behaviour recognition have

[1] Source: http://www.airforce-technology.com/projects/predator/specs.html

had limited noise, and have only evaluated overall system performance. We report the performance of flat Hidden Markov Models (HMMs) under varying degrees of noise, evaluating recognition at the lowest levels of a plan hierarchy. This indicates the challenges that face higher levels of UxV plan recognition.

A second research contribution is the application to a real UxV in a spatially unconstrained world. In visual surveillance tasks, the location based states that are traditionally employed tie recognition to the camera's perspective. In contrast, our approach provides more adaptable recognition by converting trajectories to low level actions. This allows our architecture to recognise behaviour from completely different locations and is tested against both real and simulated data. Furthermore, to the authors' knowledge this is the first work performing probabilistic behaviour recognition on a non-simulated, spatially unconstrained agent.

2 Problem Analysis

The complex domains in which UxVs operate presents a number of challenges for behaviour recognition. This includes the absence of location based states that are commonly used for visual applications (e.g. Nguyen et al., 2005). The spatial areas in which UxVs operate are vast, ranging from a hundred square metres to miles. To discretise a vehicle's location to states is impractical and could require training data for each deployment location.

Furthermore, there are a number of sources of noise that do not exist in controlled environments. This includes restrictions on sensor accuracy and frequency, in which communications failures and obstructions may prevent or delay observations. Another source of noise is the dynamic environment, where wind currents, water currents and unstable terrain can all affect the perception of actions. For example, an unmanned aerial vehicle may be blown south-west while attempting to fly south. For an observer to correctly determine the vehicles intent is must incorporate uncertainty about the vehicle's behaviour.

In any kind of surveillance domain it is reasonable to assume that agents will attempt to disguise their behaviour to prevent detection. Therefore, observations may differ from stereotypical behaviour by containing surplus actions that cause false recognition, or miss actions that have remained undetected. A simplistic approach for filtering these actions is to consider cloaking behaviour as noise.

This paper considers performing post mission analysis on an unmanned underwater vehicle. This can be used to identify the active goals of a re-planning agent in which events trigger changes in behaviour. For example, an advanced mission in sea-mine countermeasures involves processing sea-bed images while manoeuvring over an area. The detection of a suspicious object may interrupt this search behaviour to gather further data, such as chemical analysis or closer inspection. Sub-goal recognition can explain vehicle movement by identifying the active behaviours, distinguishing between searches, analysis and obstacle avoidance.

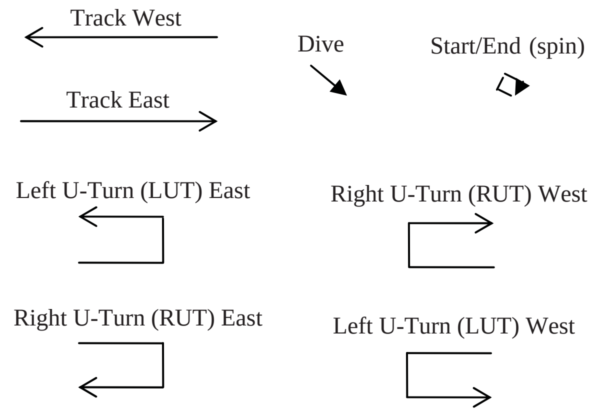

Figure 1: Common sub-goals for Unmanned Underwater Vehicles

3 Problem Formulation

Figure 1 illustrates some of the behaviours frequently executed by the underwater domain. These behaviours are achieved by performing heading based movements in the planes (x,y,z) giving sub-goal set G (1) and action set S (2).

G = {Dive, Right-U-Turn-East, Right-U-Turn-West, Left-U-Turn-East, Left-U-Turn-West, Track-East, Track-West, End} (1)

S = {Dive, Right-Turn, Left-Turn, North, North-East, East, South-East, South, South-West, West, North-West} (2)

To commence a sub-goal the vehicle first selects an action $s_i \in S$ at time step t = 0. This action is denoted as q_0 and is chosen according to the prior probability distribution π_g. For each time step the next action (q_{t+1}) is then selected according to the transition matrix A_g, in which $a_{i,j} = P(q_{t+1} = s_j \mid q_t = s_i)$. A mission is therefore constructed from sequences of sub-goals in which each element in G defines a policy $<A_g, \pi_g>$.

3.1 Vehicle Observation

Post-mission analysis is performed by obtaining vehicle longitude, latitude and depth from real UxV log files at a frequency of 1 per second. Because GPS data is unavailable underwater the vehicle approximates its position using underwater transponders. Once this has been done the vehicle continues to infer its location by using its previous position, knowledge of recent actions and Doppler based navigation. Cycling between these techniques provides a reasonable approximation of GPS position.

Our architecture uses three modules to convert changes in the vehicle's location to recognised action states. The key objective of this step is the necessity to recognise behaviour in new locations. We achieve this by borrowing techniques from online character recognition that are invariant to scale and rotation.

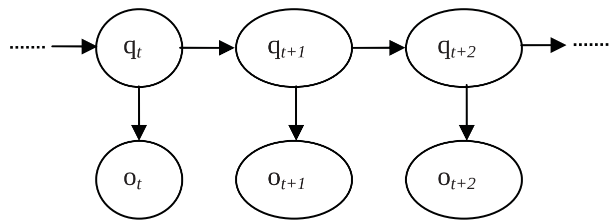

Figure 2: Directed edges assert conditional dependence in a Hidden Markov Model.

The first module calculates vehicle heading at 20 second intervals to provide one of eight states from the set {North, North-East, East,…, NorthWest}. These heading are used by the second module to detect left and right turns, as suggested by [Agarwal and Kumar, 2005]. Turns are detected as consecutive increases or decreases in direction (e.g. North, North-East, East → right turn). A unique benefit of this approach is that turns become scale independent when executed in 60 seconds or more. This lower bound is determined by the direction module, in which three changes in direction cannot be detected in less than 60 seconds.

The final module detects movement in the z-axis (depth). To overcome minor fluctuations in vehicle depth the module requires three increases of depth, each of at least one metre. The module therefore fails to emit a dive state for any observed dive of less than 3 metres. However, the module makes no constraint upon the number of observations required to detect a dive, allowing the detection of slow dives that produce many observations.

3.2 Parameter Learning

Estimating vehicle policies from observation requires action sequences in which the underlying sub-goal is both constant and known. For this purpose, we constructed a corpus of UxV sub-goal behaviours using log files from a real Hydroid Remus-100™ UxV. The corpus consisted of manually labelled partial missions in which each example demonstrated one of the eight behaviours shown in figure 1.

For a given observation sequence $O = (o_1, o_2, .., o_n)$ in which each element captures a single action, the probability distribution $P(O_t = s_i \mid q_t = s_j)$ is effected by noise. This can be modelled as an HMM by considering actions as the hidden states that emit output tokens from S. This is illustrated in figure 2 in which directed edges denote the dependence of o_{t+1} on q_{t+1}, and q_{t+1} on q_t. A model $\lambda_g \in G$ is parameterised by $<A_g, \pi_g, E_g>$, in which E_g are the emission probabilities $E_j(k) = P(O_t = s_k \mid q_t = s_j)$.

One method for estimating these parameters is to use the Baum-Welch algorithm in which the $P(O|\lambda_g)$ is maximised through iterative update and improvement. We skip the details of this algorithm but refer the interested reader to Rabiner [1989].

It is important to note that modelling the transitions between policies is beyond the scope of this paper. Instead, we focus on the ability to distinguish between the underlying sub-goals under increasing observation noise.

3.3 Recognition

To recognise sub-goal behaviour we consider the eight HMM parameters sets as classifier parameters. To classify a sequence of observations O we proceed by calculating $p(O|\lambda)$ for each model λ_g. This is performed using the forward element (α) of the Forward-Backward procedure from Rabiner [1989] and is inductively calculated in three stages:

Stage 1 – Initialisation

$$\alpha_t(i) = \pi_i \, \mathrm{E}_i(o_1) \quad 1 \le i \le |S| \tag{3}$$

Stage 2 – Induction

$$\alpha_{t+1}(j) = \left[\sum_{i=1}^{|S|} \alpha_t(i) a_{ij} \right] \mathrm{E}_j(o_t) \quad \begin{array}{l} 1 \le j \le |S| \\ 1 \le t \le T\text{-}1 \end{array} \tag{4}$$

Stage 3 – Termination

$$P(O \mid \lambda_g) = \sum_{i=1}^{|S|} \alpha_T(i) \tag{5}$$

To calculate the probability of observing $O = \{o_1, o_2, \dots, o_T\}$ the initialisation step calculates the joint probability of observing o_1 under state i. The induction step then calculates the probability of transitioning to state j at t = 2 for all states at t = 1 (a_{ij}), given observation o_2. This induction continues until t = T-1, the penultimate observation. The final probability is then obtained as the joint probability $P(o_1, o_2, \dots, o_T, q_T = i|\lambda_g)$ via the termination stage and gives $P(O|\lambda_g)$.

Model λ_g is selected as the most probable source of O iff $P(O|\lambda_g) \ge T * \max(P(O|\lambda_i)) \; \forall \; \lambda_i \ne \lambda_g \in \lambda$, where T is some threshold parameter ≥ 1. The sequence is then classified as model λ_g if it is the most probable model after K observations, or if λ_g is the most probable for C consecutive observations. We refer to this latter approach as model convergence and highlight that two or more models with the same probability will prevent convergence and result in an 'Unknown' classification.

Because convergence times vary according to sub-goal similarity, a variable length approach places less restrictions upon the number of observations required for classification. In contrast, a fixed length approach requires K (observations) = max(V), where V is the set of minimum convergence times for each sub-goal. The variable length approach should therefore out perform fixed length classification by providing earlier classification. In our main results we adopt a variable length approach, although we also assess the relative merits of both methods by comparing performance over a range of values for T, K and C.

3.4 Observation Corruption

It is commonly assumed that the elements of O have a known and constant accuracy, and that all behaviour is observed. In UxV domains the presence of sensor obstruction and availability constraints makes these assumptions invalid. These factors, and additionally cloaking actions, all result in observation sequences that incorrectly represent action transitions. For example, in the action sequence (s_1, s_2, s_3, s_4), failure to observe s_3 results in the observation sequence ($o_1 = s_1$, $o_2 = s_2$, $o_3 = s_4$). The state transition $s_2 \rightarrow s_4$ is now unrepresentative of stereotypical behaviour and will reduce the likelihood in equation 4. This effect causes a problem when multiple sub-goals share similar behaviour due to similar model probabilities. This prevents model convergence and results in an 'Unknown' classification.

To evaluate the feasibility of HMM classifiers for incomplete observation sequences, we perform sub-goal recognition upon observation test sets with 10-90% corruption. It is important to highlight that for state transition models such as HMMs, noise is defined by the number of corrupted state transitions rather than the number of states. This is because replacing the sequence {East, East, East} with {East, South-East, East} results in a 100% change in state transitions, but only a 33% change in states.

To perform observation corruption, an idealised test set was constructed through simulation in which behaviours contained zero noise. Noise was then inserted into the test set by choosing actions with uniform probability from the categories: Direction, Dive, and Turn. This noise model was selected through empirical observation of real data in which Dives and navigation errors are frequently observed. Underwater vehicles are commonly operated in altitude mode, in which the vehicle maintains a defined altitude (e.g. 3 metres) above the sea-bed. The depth of the vehicle therefore changes in line with the sea-bed terrain irrespective of the current behaviour. Section 3.1 highlighted some of the techniques used for underwater navigation due to the absence of GPS signals. A side effect of these techniques is drift in vehicle position, in which a vehicle's actual location is different from its inferred location. This results in irregular behaviour in which the vehicle attempts to correct (or fails to correct) its position. The final type of noise (Turns) simulates the presence of unobserved actions by causing state transitions that may be unlikely.

To maximise realism, direction noise was restricted to veering from the true direction. For example, if the action at time step 2 was North (A(2)=North), direction noise at time step 2 could only be A(2)=North-West or A(2)=North-East, realistically simulating drift. Furthermore, in the occurrence of two continuous direction noise states, the second insertion is restricted by the first to simulate increasing drift. For example, if noise A(2)=North-West, A(3) may only be West or North. We refer to this kind of noise as contiguous.

4 Related Work

Plan recognition has generally followed two approaches: logic based methods, such as Kautz [1991], or probabilistic methods, such as [Blaylock and Allen, 2006]. For the UxV domain, environment noise makes complete world knowledge impossible to attain, breaking the assumption of logic based methods. In contrast, probabilistic approaches such as the Hidden Markov Model consider incomplete knowledge as uncertainty, allowing inference on previously unseen action sequences.

The Abstract Hidden Markov Model (AHMM) is a variation on the standard structure and is applied to a visual surveillance task by [Bui and Venkatesh, 2002]. They track agent trajectories using video in a complex environment of corridors and rooms. [Nguyen et al., 2003] extend the AHMM to allow more complex state dependencies in a similar environment, naming it the Abstract Hidden Markov mEmory Model (AHMEM). In both approaches states are obtained by segmenting the environment into regions, emitting state observations as each is traversed. This causes a dependency between model parameters and the environment's locations, enforcing model retraining for a new environment. Our work differs from Bui and Nguyen by recognising sequences of actions rather than locations.

Higher level behaviours are frequently composed of reusable components, such as walking from the cupboard to the fridge while making either a snack or a meal. To construct robust and scalable recognition the number of model parameters must be restricted by sharing reusable behaviours [Nguyen et al., 2005]. Nguyen and colleagues demonstrate that for scenarios such as these the Hierarchical Hidden Markov Model (HHMM) offers superior performance over a flat HMM. Like [Bui and Venkatesh, 2002] the state space is again determined by segmenting the environment, requiring parameter relearning to detect behaviour within a different environment.

In some ways our approach is quite similar to the inference of transportation routines by [Liao et al., 2007a]. They use sequences of Global Positioning System (GPS) coordinates to predict a person's destination location (goal), transportation mode (e.g. foot) and future movement (e.g. turns left at intersection). To robustly handle GPS coordinates they 'snap' locations to the nearest known street to provide some degree of adaptability for variations in behaviours.

In our own work we also use GPS observations, although these are inherently noisier than in [Liao et al., 2007a]. Furthermore, we translate sequences of GPS coordinates into simple headings to achieve location independence. This is not true of [Liao et al., 2007a], in which a new city would require the construction of a new location grid.

[Liao et al., 2007b] use hierarchical Conditional Random Fields (CRFs) to extract activities and important locations from GPS data. CRFs are a discriminative technique and are less restrictive than generative based methods. Generative techniques (e.g. HMMs) assume observations are independent given the state, while discriminative techniques make no such assumptions. As with other related work they perform recognition on known environments and cannot adapt to new locations. Furthermore, they cite a drawback of discriminative learning as requiring more training data than generative approaches [Ng and Jordan, 2002]. The practical implications of running vehicle missions restricts the availability of UxV training data and in the short term may prevent the use of CRFs.

Our work also shares similarities to [Sukthankar and Sycara, 2006] in which they use flat HMMs to distinguish between 3 spatially invariant behaviours. The main differences from our own work is that their observations are obtained from a simulated world and that Gaussian noise is inserted at the (x, y) coordinate level. In our work the training corpus is obtained by observing a real unmanned vehicle and the inserted noise in fundamentally different. It is unclear what states are derived from the noisy coordinates in [Sukthankar and Sycara, 2006], however, our states represent 3 to 20 second intervals and it is at this level that noise is added.

5 Results

To explain the metrics gathered during the study they will be described in terms of an example. The performance of the 'Track-West' model was attained by presenting 'Track-West' samples to all eight models. These samples were provided in a sequential manner, increasing their length by one action each iteration. The probability of a model generating the sequence can be calculated and compared to other models, allowing the most probable model to be selected.

Convergence is attained when a model remains at least twice as probable as any other for three consecutive iterations. A Track-West *Convergence Recall* of 0.9 would therefore indicate that for 90% of the Track-West test samples, the Track-West model converged as the most probable model.

Convergence Precision complements this metric by calculating the accuracy of predictions. For example, the total number of times the 'Track West' model converged correctly across all test sets divided by the number of times the 'Track West' model converged.

These two metrics measure how often a sequence type is correctly identified and how accurate each classifier (model) is. To assess recognition performance under 'standard' conditions, table 1 presents the average accuracy and recall of all 8 models after performing 3-fold cross validation on the real UxV corpus. This represents the scenario in which we assume full observation of a non-cloaking vehicle. Under these conditions, both precision and recall remain very high at 0.99 and 0.93.

However, of primary interest is the performance of the classifiers under additional levels of noise, mimicking missed observations (through obstruction/communications failure) and cloaking behaviour (adversarial agents). For these scenarios, the models were retrained using the entire real UxV training corpus, and then tested with simulated data. Figure 3a shows the convergence recall for each of the learnt models, in which recall remains above 0.85 at noise levels up to 30%. At higher noise levels recall is lower, although remaining above 0.7 at 50% noise. It is interesting to note that the Dive model has the highest level of recall because it consists largely of one action, making confusion with other models less likely.

Poor performance at higher noise levels can be explained by considering the length of action sequences. A 'Right U-Turn East' sequence such as (6) could be transformed into (7) by 70% corruption.

{East, East, East, South-East, **Right-Turn**, South, **South**, South-West, **Right-Turn**, West} (6)

{East, East, East, South-East, **East**, South, **South-East**, South-West, **West**, West} (7)

Several high probability state transitions have now been eliminated from the sequence, an affect that would be less likely to occur in longer sequences. The net effect of this loss in key state transitions is a higher level of confusion between models. To gain a better understanding of this confusion table 2 presents the confusion matrix for 70% noise. Each row presents the distribution of classifications (convergence) for a single type of sequence. For example, the Dive row shows that 6% of Dive sequences were classified (or confused) as 'Right U-Turn East' sequences. If no model converged as the most likely the sequence remains unclassified, resulting in some rows adding up to less than 100%.

The matrix illustrates that the confusion between some models is highly logical. For example, 'Track West' sequences are classified as 'RUT West' 16% of the time. Analysing the noise shows that at a level of 70% it is highly likely that some of the noise is contiguous (e.g. North-West → North). This leads to state transitions such as (8) which are very similar to the 'Right U-Turn West' behaviour (9).

{*West*→*West*→*North-West*→*North*} (8)

{*West*→*West*→*North-West*→*North*→North-East→East}(9)

Similar confusion occurs between several of the models, however all of the models are confused to some degree with the 'End' behaviour. The reason for this lies in the nature of the 'End' activity in which the vehicle floats on the surface. The motion of waves causes the vehicle to spin which leads to high transition probabilities for any direction state to an adjacent one. For example, a vehicle spinning right is likely to move from North to North-East, while a vehicle spinning left is likely to move from North to North-West. The number of adjacent state transitions in any sequence will increase with direction noise, causing confusion with the 'End' model.

Test Set	**Avg. Precision (across 8 models)**	**Recall (across 8 models)**
Set 1	0.98	0.94
Set 2	0.99	0.89
Set 3	0.98	0.97
3-Fold Average	***0.99***	***0.93***

Table 1: Average precision and recall using 3-fold cross validation on real UxV Corpus

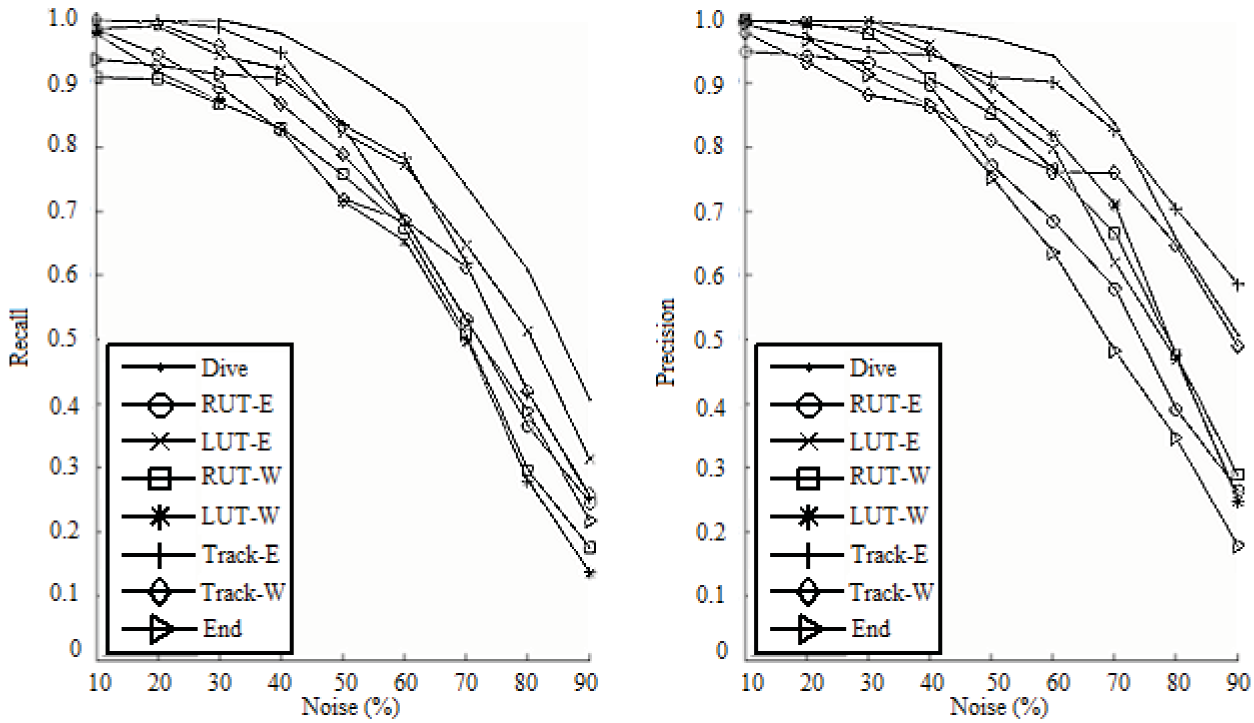

Figure 3: (a – Left) Sub-Goal Recall vs. Noise and (b – Right) Classifier Precision vs. Noise

Sequence\Model	**Dive**	**RUT East**	**LUT East**	**RUT West**	**LUT West**	**Track East**	**Track West**	**End**
Dive	0.74	0.06	0.10	0.03	0.04	0.02	0.00	**0.01**
RUT East	0.01	0.70	0.04	0.01	0.01	**0.06**	0.02	**0.14**
LUT East	0.01	0.06	0.77	0.00	0.01	**0.06**	0.00	**0.07**
RUT West	0.05	0.04	0.04	0.62	0.05	0.00	**0.10**	**0.09**
LUT West	0.01	0.02	0.03	0.03	0.66	0.00	**0.20**	**0.04**
Track East	0.03	**0.08**	**0.10**	0.01	0.01	0.59	0.00	**0.17**
Track West	0.05	0.02	0.04	**0.16**	**0.11**	0.00	0.58	**0.03**
End	0.08	0.11	0.11	0.06	0.07	0.01	0.01	0.56

Table 2: The confusion Matrix for sub-goal classification under 70% noise.

Param. Values	**Avg. Precision**	**Avg. Recall**
C=4, T=4	0.59	0.43
C=3, T=3	0.55	0.47
C=4, T=3	0.59	0.44
C=3, T=4	0.57	0.48
K=8	0.55	0.54
K=10	0.59	0.57

Table 3: The affect of parameters C (minimum converged observations), T (convergence ratio) and K (fixed observation length) at 70% noise.

Looking next at the convergence precision (figure 3b) it is clear that a high precision rate is achieved at low noise levels. In general, precision remains above 0.78 at up to 50% noise, but falls linearly after this point. It is interesting to observe that at 70% noise and above there are three models that perform better than the others. This separation is most apparent at 90% noise where the 'Dive', 'Track-West' and 'Track-East' models outperform the others by at least a factor of 2. This is because each of these three sub-goals consists largely of one state and therefore lack 'key transitions'. As previously discussed, the elimination of key transitions is one of the reasons for confusion between models. These models however have fewer high probability transitions making them less susceptible to noise.

Noise	Dive	RUT East	LUT East	Track East	End
10%	0.92	0.05	0.03	0.00	0.00
20%	0.80	0.10	0.09	0.00	0.00
30%	0.68	0.14	0.17	0.00	0.01
40%	0.60	0.18	0.18	0.00	0.03
50%	0.56	0.20	0.21	0.00	0.03
60%	0.48	0.20	0.25	0.00	0.08
70%	0.45	0.22	0.26	0.00	0.07
80%	0.36	0.23	0.30	0.00	0.12
90%	0.33	0.27	0.25	0.00	0.14

Table 4: Confusion matrix for the '*Track West*' test set when only 5 classifiers are present

An unfortunate side affect of the UxV domain is the lack of previous work with which results may be compared. While similar techniques have been employed in [Sukthankar and Sycara, 2006], there is no clear mapping between states to allow a subjective comparison of performance. However, one aspect that can be considered is the relative merit of variable length classification over (traditional) fixed length. Recall that in variable length classification, a sequence O is classified as λg if:

$$P(O|\lambda_g) \geq T * \max(P(O|\lambda_i)) \ \forall \ \lambda_i \neq \lambda_g \in \lambda$$

for C consecutive iterations. In contrast, fixed length classification classifies the partial observation sequence $\{o_1, .., o_k\}$ as λg if $P(o_1, .., o_k|\lambda_g) \geq \max(P(o_1,\ldots,o_k|\lambda_i)) \ \forall \ \lambda_i \in \lambda$. In table 3 the average precision and recall across all eight models is shown. In these scenarios performance was obtained for different variable length parameter values (C and T) and fixed lengths (K) under 70% noise. A fixed length of 10 offers best overall performance, although similar precision may be obtained with variable length classification at the detriment of recall. Applications requiring early classification, should as on-line behaviour recognition, should therefore consider the variable length approach. However, for less offline applications the affect on recall suggests that superior performance is achieved by deriving the optimal value of K.

Assigning values for (C, T) or K requires arbitrary initialisation and may be further tuned using optimisation techniques such as Gradient Descent. In this paper, initialisation values were discovered by hand within a reasonable time and were not further optimised. Table 3 illustrates that the change in precision and recall are small between differing parameter assignments.

Lastly, we briefly discuss the affect of removing 'RUT West', 'LUT West' and 'Track West' models from the classifier set. Presenting all test data to the remaining models then mimics an adversary scenario in which not all behaviours are known. Table 4 shows that at low noise levels the 'Track West' test sequences were most frequently classified as Dives. This is because the 'Dive' sub-goal contains both horizontal and vertical motion, allowing model convergence. In some ways, this classification may be viewed as positive, as the most similar behaviour to the observations was still correctly identified. However, as the noise increases the confusion spreads across the different models, and while not shown in table 4, convergence frequency remains high throughout. This is an undesirable effect as it shows that completely unknown behaviours may still be classified.

6 Conclusion and Future Work

We have applied behaviour recognition to the new domain of unmanned vehicles (UxVs). Previous work has performed behaviour recognition on human agents but their solutions have largely been constrained to known environments. This prevents a trained model from recognising behaviours in a new location without at least some degree of retraining. Our work has overcome this barrier by using a location independent state representation and has achieved 0.99 recognition precision and 0.93 recall on real UxV data.

We have identified several sources of noise that make behaviour recognition particularly challenging for UxV applications. Although previous research has utilised noisy sensor readings, location based states restrict confusion potential. In our location independent architecture this potential is much greater, with many behaviours indistinguishable from others at high levels of noise. To obtain a firm understanding of the affect of noise for this domain we performed behaviour recognition with flat hidden Markov models under varying levels of noise. Classification precision remained high at noise levels up to 50%.

These results are encouraging for noise levels below 50%, although the significant drop in convergence precision at higher noise levels may have implications. In hierarchical recognition the detected sub-goals are used to recognise the parent goal. These goals are likely to have few state transitions, making the impact of noise more profound. For example, assume goal 1 has the sub-goal sequence A→B→C, and that sub-goal B is incorrectly classified as sub-goal D. The observed sequence now has no transitions in common with the original goal. This aspect may therefore cause difficulty in recognising plans with few sub-goals, and will be investigated in further work by performing hierarchical recognition.

In our current work sub-goals are direction dependent, requiring multiple models to recognise different orientations of the same behaviour. A natural extension to this work would be the application of a scale and rotation invariant technique that would facilitate fewer sub-goal models. At the same time, one method by which recall may be improved is to reduce the number of key transitions, an aspect of the current state representation that has been shown to reduce recall performance at high levels of noise.

A further extension of this work is to perform behaviour recognition on a vehicle's observations of a second vehicle. GPS coordinates would encapsulate the vehicles own uncertainty about its location, and would also encapsulate uncertainty about the vehicle being observed.

To further address scenarios in which some behaviours are unknown, it may be interesting to combine model convergence with a minimum probability threshold. This technique was suggested by Laskey [1991] as a method of identifying model mismatch and may be useful in preventing the incorrect classification of unknown behaviours with very low probabilities.

Acknowledgments

This work is funded by the UK Ministry of Defence under the Competition of Ideas.

References

Agarwal, S. and Kumar, V. 2005. Online Character Recognition. In Proceedings of the Third International Conference on Information Technology and Applications (ICITA'05). 2:698-703

Blaylock, N., and Allen, J. 2006. Hierarchical instantiated goal recognition. In AAAI Workshop on Modeling Others from Observations (MOO-2006). AAAI Technical Report WS-06-13. Pages 8-15

Bui, H. H., and Venkatesh, S. 2002. Policy recognition in the Abstract Hidden Markov Model. *Journal of Artificial Intelligence Research*. 17: 451-499.

Kautz, H. A. 1991. A Formal Theory of Plan Recognition and its Implementation . *Reasoning About Plans*. Allen, Kautz, H.A., Pelavin, R.N., and Tenenberg, J.D. Eds. Morgan Kaufmann Publishers. Pages 69-126

Laskey, K.B. 1991. Conflict and Surprise: Heuristics for Model Revision. In Proceedings of the 7th Conference on Uncertainty in Artificial Intelligence. Pages 197-204

Liao, L., Patterson, D.J., Fox, D., and Kautz, H. 2007a. Learning and inferring transportation routines. *Artificial Intelligence*. 171(5-6):311-331

Liao, L., Fox, D., and Kautz, H. 2007b. Extracting Places and Activities from GPS Traces Using Hierarchical Conditional Random Fields. *International Journal of Robotics Research*. 26(1):119-134

Nau, D., Ilghami, O., Kuter, U., Murdock, J. W., Wu, D., and Yaman, F. 2003. SHOP2: An HTN planning system. *Journal of Artificial Intelligence Research*. 20:379-404

Nguyen, N.T., Bui, H. H., Venkatesh S., and West. G. 2003. Recognising and monitoring high-level behaviours in complex spatial environments. In Proceedings of IEEE International Conference on Computer Vision and Pattern Recognition. Pages 620-625.

Nguyen, N.T., Phung, D.Q., Venkatesh, S., and Bui, H. 2005. Learning and Detecting Activities from Movement Trajectories Using the Hierarchical Hidden Markov Model. In Proceedings of The 2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05). 2: 955-960

Oliver, N., Garg, A., and Horvitz, E. 2004. Layered representations for learning and inferring office activity from multiple sensory channels. *Computer Vision and Image Understanding*. 96(2):163-180

Rabiner, L. R. 1989. A tutorial on hidden markov models and selected applications in speech recognition. *Proceedings of the IEEE*, 77(2):257-286

Sotzing, C. C., Evans, J., and Lane, D.M. 2007. A Multi-Agent Architecture to Increase Coordination Efficiency in Multi-AUV Operations. *IEEE Oceans Europe 2007*.

Sukthankar, G. and Sycara, K. 2006. Robust Recognition of Physical Team Behaviors Using Spatio-Temporal Models. In Proceedings of Fifth International Joint Conference on Autonomous Agents and Multi-Agent Systems. Pages 638-645