Bridging the Sense-Reasoning Gap using DyKnow: A Knowledge Processing Middleware Framework

Fredrik Heintz, Piotr Rudol, and Patrick Doherty

Department of Computer and Information Science
Linköpings universitet, Sweden
{frehe, pioru, patdo}@ida.liu.se

Abstract. To achieve complex missions an autonomous unmanned aerial vehicle (UAV) operating in dynamic environments must have and maintain situational awareness. This can be achieved by continually gathering information from many sources, selecting the relevant information for current tasks, and deriving models about the environment and the UAV itself. It is often the case models suitable for traditional control, are not sufficient for deliberation. The need for more abstract models creates a sense-reasoning gap. This paper presents DyKnow, a knowledge processing middleware framework, and shows how it supports bridging the gap in a concrete UAV traffic monitoring application. In the example, sequences of color and thermal images are used to construct and maintain qualitative object structures. They model the parts of the environment necessary to recognize traffic behavior of tracked vehicles in real-time. The system has been implemented and tested in simulation and on data collected during flight tests. [1]

1 Introduction

Unmanned aerial vehicles (UAVs) are becoming commonplace in both civil and military applications, especially for missions which are considered dull, dirty and dangerous. One important application domain for UAVs is surveillance. Such missions may involve flying over unknown areas to build terrain models, to quickly get an overview of a disaster area including helping the rescue services to find injured people and deliver medical supplies, or to help law enforcement agencies to monitor areas or people for ongoing or potential criminal activity. To achieve these complex missions an autonomous UAV must continuously gather information from many different sources, including sensors, databases, other UAVs, and human operators. It then selects relevant information for the current task, and derives higher-level knowledge about the environment and the UAV itself in order to understand what is happening and to make appropriate decisions. In other words, the UAV must create and maintain its own situational awareness in a timely manner.

To create situation awareness a UAV needs to build models of the environment and use them to reason contextually about the past, current, and future state of the world.

[1] This work is supported in part by the National Aeronautics Research Program NFFP04 S4203 and the Strategic Research Center MOVIII, funded by the Swedish Foundation for Strategic Research, SSF.

These models should be constructed from information gathered from distributed sources and aggregated in a timely manner in order to capture the latest developments. Since there are many models that could be built and since a UAV has limited resources it is important that the appropriate models are constructed contextually for the particular task at hand. When the task changes this should be reflected in the models as well.

What is an appropriate model will depend on what properties of the world are relevant, what reasoning is needed to make decisions to achieve a task, and the context within which that reasoning is made. One functionality is the construction of models from data aggregated from different sensors that can be used to reason about the environment and the UAV in real-time. There are numerous approaches to building quantitative models based on sensor data. These models are suitable for traditional tracking and control applications but do not provide appropriate abstractions when reasoning about complex situations such as traffic. On the other hand, there are many qualitative modeling approaches using formal symbols which are well suited to do high level reasoning about the environment. How to connect these different approaches and to close the gap between sensing and reasoning is still an open research question.

This paper presents an implemented traffic monitoring application that uses the knowledge processing middleware framework DyKnow [1,2] to bridge the sense-reasoning gap. It is done by creating tailored models at different levels of abstraction as described by declarative policies. The models are interconnected in order to describe dependencies and to keep them updated. The models created can be used to reason qualitatively about the world as it develops using for example temporal logic and a complex event formalism called chronicle recognition.

2 Traffic Monitoring

Imagine a human operator trying to maintain situational awareness about a traffic situation in an urban area using UAVs which look for accidents, reckless driving, or other relevant activities. One approach would be for one or more UAVs to relay videos and other data to an operator for human inspection. Another more scalable approach would be for the UAVs to monitor traffic situations which arise and only report back the high level events observed. This would reduce the amount of information generated and help an operator focus attention on salient events. This paper describes such a traffic monitoring application where cars are tracked by a UAV platform and streams of observations are fused with a model of the road system in order to draw conclusions about the behavior of cars.

The input consists of images taken by color and thermal cameras on a UAV which are fused and geolocated into a single world position. This stream of positions is then correlated with geographical information system (GIS) data in order to know where in a road system an object is located. Based on this information, high level behaviors such as turning at intersections and overtaking are recognized in real-time as they develop using a chronicle recognition system.

An overview of the components of the traffic monitoring application is shown in Fig. 1. The three sensors used, the two cameras and the helicopter state estimation (which is fused from inertial and GPS data), are shown to the left. These provide the

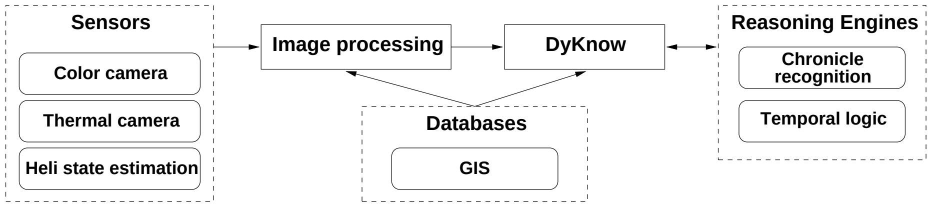

Fig. 1. Overview of the components of the traffic monitoring application.

low level data about the environment and the UAV. The next component is an image processing system which tracks objects seen by the cameras. When an object is being tracked, images from the two cameras are fused to provide an estimation of the position in the world of the tracked object. Each time a new frame is analysed a new position estimate is produced. From this stream of position estimations DyKnow derives further abstractions used to recognize high level events and closes the sense-reasoning gap.

To describe a high-level event a formal representation called a *chronicle* is used [3]. A chronicle defines a class of events using a simple temporal network (STN) [4] where the nodes are primitive events and the edges are temporal constraints between event occurrences. The chronicle recognition engine takes a stream of primitive events and detects all chronicle instances. An instance is detected if the stream contains a set of event occurrences which satisfy all the constraints in a chronicle model. The chronicles used in this application contain primitive events which capture the structure of the road network, qualitative information about cars such as which road segment they are on, and qualitative spatial relations between cars such as beside and behind. Creating a stream of events based on sensor data which accurately represents the environment of the UAV, is a concrete instance of the sense-reasoning gap.

To bridge the gap, DyKnow takes a stream of position observations provided by the image processing system and derives an event stream representation of cars and qualitative spatial relations between cars. DyKnow also derives an event stream representation of the road network from the information stored in the GIS. One issue that must be handled is how to anchor car symbols used in the chronicles to objects being tracked [5]. Since the image processing system may lose track of cars or start tracking other non-car objects, DyKnow has to dynamically estimate and continually monitor the type and identity of objects being tracked. To do this, the normative behavior of different objects and the conditions for assuming that two objects have the same identity are described using temporal logic. When a tracked object is found which satisfies the normative behavior of e.g. a car, a new car representation is created and the tracked object is *linked* to the new car representation. From this moment the car representation will be updated each time the tracked object is updated. Since links only represent hypotheses, i.e. they are always subject to becoming invalid given additional observations, the UAV continually has to verify the validity of the links. This is done by monitoring that the normative behavior of the assumed object type is not violated. For example, an object assumed to be a car must not violate the normative constraints on cars, e.g. leaving the road. If it does violate the constraint, then the corresponding link is removed, in other words the object is no longer assumed to be a car.

To evaluate a temporal logical formula, DyKnow has to derive a model representing the value over time of the variables used in the formula. These values must be synchronized in time, so that the evaluation mechanism receives a state for each time-point containing the value of each of the variables at that time-point. This is done by defining a policy for each of the formulas which DyKnow uses to derive the required model. Since these models are derived from sensor data, it is another concrete example of how DyKnow can be used to bridge the sense-reasoning gap. The evaluation is done using progression which means that the evaluation is performed in real-time as soon as a new state is available. This means that the truth value of a formula will be derived as soon as it is possible.

The application has been tested both on simulated cars driving in a road system and on real data captured during flight tests.

3 Conclusions

A traffic monitoring application which is an instance of a general approach to creating high-level situation awareness has been presented. The implemented system takes as input sequences of color and thermal images. They are used to construct and maintain qualitative object structures and recognize the traffic behavior of the tracked vehicles in real-time. The system is tested both in simulation and on data collected during flight tests. We believe that this type of system where streams of data are generated at many levels of abstraction using both top-down and bottom-up reasoning handles many of the problematic issues related to closing the sense-reasoning gap in robotic systems. One reason for this is that the information derived at each level is available for inspection and use at all times. This means that the subsystems have access to the appropriate abstraction while it is being continually updated with new information and used to derived even more abstract structures. High-level information, such as the type of vehicle, can then be used to constrain and refine the processing of lower level information. The result is a very powerful and flexible system capable of achieving and maintaining high-level situation awareness.

References

1. Heintz, F., Doherty, P.: DyKnow: An approach to middleware for knowledge processing. Journal of Intelligent and Fuzzy Systems **15**(1) (2004) 3–13
2. Heintz, F., Doherty, P.: A knowledge processing middleware framework and its relation to the JDL data fusion model. Journal of Intelligent and Fuzzy Systems **17**(4) (2006) 335–351
3. Ghallab, M.: On chronicles: Representation, on-line recognition and learning. In: Proceedings of the Fifth Intl Conf on Principles of Knowledge Representation and Reasoning. (1996)
4. Dechter, R., Meiri, I., Pearl, J.: Temporal constraint networks. AIJ **49** (1991)
5. Coradeschi, S., Saffiotti, A.: An introduction to the anchoring problem. Robotics and Autonomous Systems **43**(2-3) (2003) 85–96