A Probabilistic Network Approach to Prioritising Sensing and Reasoning: a step towards satisficing modelling

Sharon Wood
School of Cognitive & Computing Sciences
University of Sussex
Brighton, BN1 9QH, UK
email: sharonw@cogs.susx.ac.uk

Abstract

The work described here concerns situational modelling in the support of dynamic decision making when planning in the driving domain. When resources are constrained, sensing and reasoning can be dynamically prioritised by evaluating which part of an agent's situational model most needs updating to support the agent's activities. A probabilistic network approach is used to competatively prioritise modelling requirements.

Motivating Situational Modelling

Dynamic, multi-agent environments impose on the agent the need to respond in a timely manner to ongoing events. The problem of how to generate such timely responses has given rise to two modes of approach; the reactive (*cf* Agre and Chapman, 1987) or reflex architecture, and the hybrid architecture embodying more deliberative techniques (*cf* Sanborn and Hendler, 1988). Other approaches attempt to make the goal driven approach more robust under rapidly changing conditions (Georgeff and Lansky, 1987; Wood, 1993).

An intrinsic feature of many reactive systems, or those including a reactive component, is their reliance on the information held in the current situation - the world as it appears *now* - in informing an appropriate response. And yet, at any given moment, relevant information may not be entirely accessible and the response made is likely to be impoverished as a result (Wood, 1995). Constructing an enduring model of the situation, based on information accumulated from previous observations helps to guard against this pitfall (Wood, 1993). A timely response may also require the agent's *active anticipation* in predicting events to which it must react, in order to *initiate* a reaction in good time (Wood, 1993; 1995).

Situational modelling can therefore play a twofold role. It can provide a more complete picture of events than can be sensed at any given moment in time. It can also provide a picture of anticipated future events that can be used to inform the actions of the agent, thus enabling those actions to be initiated in a timely manner.

The Modelling and Sensing Relationship

Modelling situations and their outcomes in order to support timely interactions is itself subject to the same constraints as deliberating about how to respond. Limitations on sensing have led to the use of active vision techniques (Aloimonos *et al*, 1987; Bajcsy and Allen, 1984; Reece and Shafer, 1995; Tsotsos, 1992) to guide sensing in a task-related manner. This allows the agent's informational *needs* to be taken into account when prioritising what to sense.

The role of situational modelling is to serve the agent's informational needs; to the extent that it is doing this effectively, sensing and modelling complement each other. The agent's task needs might be said to be the determinant of the model's content, whilst the agent's situational modelling needs, which indirectly serve its tasks, might be said to be the determinant of what the agent senses. It is this model of the relationship between sensing and modelling that forms the framework for the work described here.

Situational modelling needs arise when the information being modelled becomes out of date. We can notionally attribute this two factors, although they amount to one and the same thing: uncertainty. There is inherent uncertainty in making predictions because we cannot necessarily identify all the factors causally related in the outcome of a situation (this is especially true in determining the behaviour of other agents (*cf* Schmidt *et al*, 1978; Wood, 1993). An agent may therefore believe an event is going to occur, but it may not, and something else may take its place. Predictions must therefore be monitored for confirmatory and disconfirmatory evidence (*cf* Huber, 1993). Where the agent's beliefs are based on direct observation, these beliefs are nonetheless subject to persistence effects (*cf*

Dean and McDermott, 1987) which sooner or later render them invalid.

Focus of Attention in Situational Modelling

The work described here focusses on the use of situational modelling to support interaction with dynamic, rapidly changing driving situations. It extends earlier work (Wood, 1993) by supplementing sensing and modelling activities with mechanisms to direct the focus of attention according to modelling needs. These needs will sometimes be in conflict with each other, and work in progress involves the application of probabilistic network techniques in evaluating and prioritising sensing requirements. The aim of the work is to enable the building of situational models that are *satisficing* (Simon, 1981) to the agent's needs whilst reducing the computational overhead involved in maintaining a complete and accurate model of the enduring situation.

The process of prioritising sensing requirements takes place in two stages, corresponding to a hierarchical decomposition of the allocating focus of attention task. The first stage corresponds approximately to determining direction of gaze in a human driver[1]. The outcome is a decision, determined probabilistically, about whether, for example, to direct gaze directly ahead, or behind, or towards a number of other locations manifest within the driving scene, for example, intersections (junctions). The topology of the road network may require further differentiation into specific carriageways, corresponding to the different directions of travel available to traffic, and into lanes on wider roads[2]. The decision combines a fixed weighting reflecting the intrinsic value to the agent of directing gaze in a particular direction, with evidence concerning the reliability of data modelled.

A probabilistic network is initialised for each driver-viewable region. The network is a singly-connected directed acyclic graph (DAG). For example, in the scene in Figure 1, regions include the give-way zones for turning left(A) and right(B); exit zones for turning left into the nearside(C) and offside(D) lanes; a turning right zone(E); an exit zone for impending traffic from the right(F); impending right and left nearside and offside zones(G,H,M,N) and a similar set of zones for retreating traffic(I,J,K,L).

[1]It could also correspond to determining which camera image to process from a collection of fixed cameras mounted in differing orientations.

[2]Note that the defined road regions, *eg* the give-way zone to an intersection, are independent of the viewer (Wood, 1993) but the choice of which region to view next would then be interpreted as a decision to, say, look straight ahead, etc, according to the relationship of the region to the observer-driver.

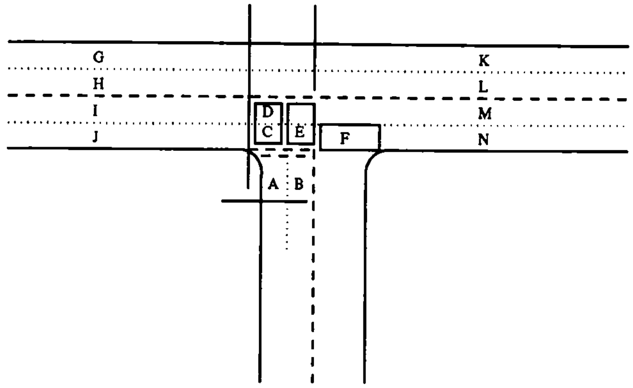

Figure 1: Example of Typical Zones at Intersection

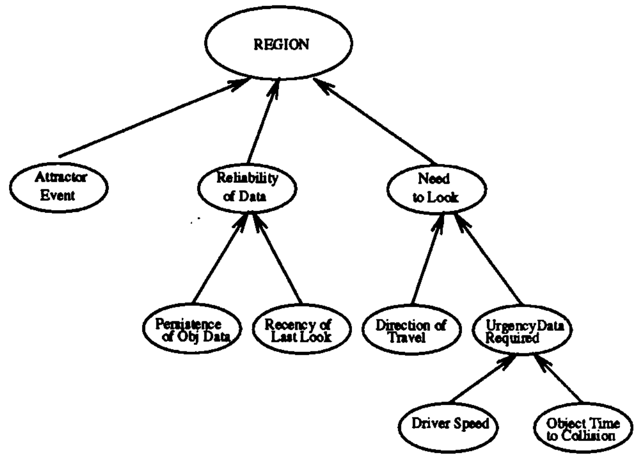

Figure 2: Network for Regions

The network is shown in Figure 2. The node, POD (Persistence of Object Data) represents the probabilistic expectation regarding the persistence of data integrity for the region. The node, RLL (Recency of Last Look) represents the probabilistic expectation that the region has been recently observed. These probabilities combine with evidence for types of (previously observed) objects in the region (not shown in diagram). This in turn combines with the expected reliability of the modelled data, prepresented by node RD (Reliability of Data). Node AE (Attractor Event) represents the probabilistic expectation that an attention-worthy event has taken place in the region.

Node DT (Direction of Travel) represents the probabilistic expectation that the region lies in the driver's pathway. The node UDR (Urgency Data Required) represents the probabilistic expectation concerning the urgency with which the region should be attended

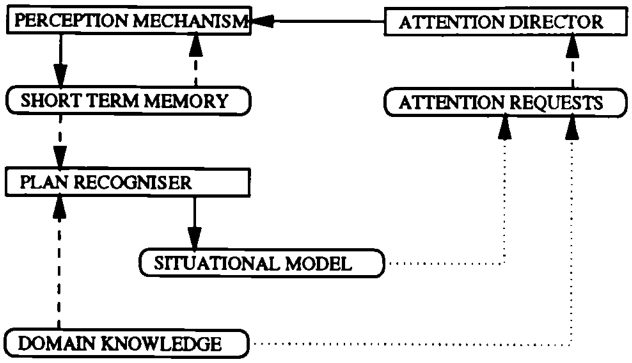

Figure 3: AutoAttend

to. Nodes DS (Driver Speed) and OTC (Object Time to Collision/Contact) represent probabilistic expectations regarding of driver and object behaviour. The node NL (Need to Look) represents the probabilistic expectation of the driver's need to observe the region. The node Region represents the probabilistic expectation that the region should be observed.

A net is initialised for each region viewable by the driver[3]. The status of the network is based on data held in the situational model. As observations are made by the driver, and the status of simulated data alters over time, the networks will be updated to provide an evaluation of the priority with which each region should be viewed.

The second stage of prioritising sensing requirements corresponds to identifying a fixational target. This will be some object, such as the vehicle in front, or location, such as the edge of the carriageway, that falls within the field of view as determined by the direction of gaze. The choice made combines a weighting reflecting the observation-worthiness of the object with evidence concerning the reliability of data currently modelled. This evidence might be viewed as a rating of belief in the persistence of belief concerning the modelled object. The rating of observation-worthiness might itself comprise several factors reflecting properties of the object, such as its type, its speed of movement[4] and its trajectory, for example.

The AutoAttend mechanism is currently undergoing integration into the AutoDrive architecture (Wood, 1993). The relevant components of this architecture (*ie* omitting the world simulation, planning and decision-making components) are shown in Figure 3. The dotted lines describe the new role of *Domain Knowledge*

[3]Regions are designated by a simulation program which generates descriptions of the driver's world for each cloned driver agent (Wood, 1993).

[4]Even fixed objects "move" relative to the viewer!

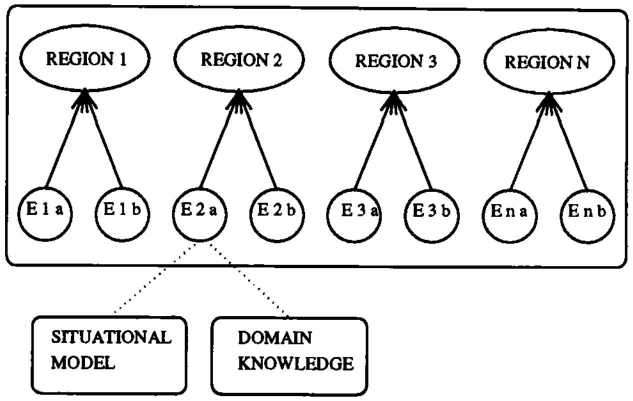

Figure 4: Attention Requests - collection of networks

on the interpretation of evidence, for example, that a region in the driver's pathway has an impact on their need to look (NL) at that region; and the new role of the *Situational Model* in providing evidence for the networks [5].

Attention Requests take the form of probabilistically determined evaluations of the need to view each of the regions within sight of the driver (see Figure 4). The role of the *Attention Director* is to select (oneof) the highest priority regions for attention by the *Perception Mechanism*[6].

Once a region has been selected, potential objects to fixate (and process further) from within the chosen region are treated in a similar manner as for regions. The objects, or *Attention Requests*, in the form of probabilistic networks[7] are evaluated and a selection made by the *Attention Director* for attention by the *Perception Mechanism.*

Discussion

The probabilistic model is currently under development, but appears to offer a means of addressing some of the problems posed by modelling under resource constraints and the construction of satisficing models.

The two stage process described above allows the monitoring of the current situation in a manner that supports the construction and maintenance of a satisficing model of the agent's scenario. Where resource limitations constrain sensing and modelling activities, the probabilistic model identifies how these would be

[5]Solid lines indicate flow of control/data generation, dashed lines indicate data flow.

[6]Although in practice this might involve a notional head movement, attention is not mediated via the agent actions in this implementation. Instead, the agent simply "ignores" data for other regions by partitioning the data available.

[7]The networks for objects are currently under construction.

most usefully targetted and efficacious in satisfying the informational needs of the agent. Where the tasks of the agent, including the task of identifying the intentions of other agents (*cf* Wood 1993; Huber, 1993), are supported by the model, they have no direct impact on this process; indirectly, they may impose a stricter requirement on the reliability of the data modelled which may be reflected in the competitative prioritisation of where to direct gaze and which object to fixate.

References

Agre, PE & Chapman, D (1987) Pengi: an implementation of a theory of activity In *Proceedings AAAI-6 Vol. 1,* Seattle, Washington, pp268-272

Aloimonos, Y, Weiss,I & Bandopadhay, A (1987) Active Vision, In *Proceedings IEEE First International Conference on Computer Vision*, London, pp35-54

Bajcsy, R & Allen, P (1991) Sensing Strategies, In *Proceedings US-France Robotics Workshop*

Dean, TL & McDermott, DV (1987) Temporal Data Base Management *Artificial Intelligence, Vol 32, pp1-55*

Georgeff, MP & Lansky, AL (1987) Reactive Reasoning and Planning In *Proceedings of AAAI-87*, pp677-682

Huber, MJ (1993) Observational Uncertainty in Plan Recognition Among Interacting Robots in *Proceedings of the Workshop on Dynamically Interacting Robots, IJCAI-93*, Lille, France.

Reece, DA & Shafer, SA (1995) Control of Perceptual Attention in Robot Driving *Artificial Intelligence, Vol 78(1-2), pp397-430*

Sanborn, J & Hendler, J (1988) A Model of Reaction for Planning in Dynamic Environments *AI in Engineering, Vol 3(2), pp95-102*

Schmidt, CF, Sridharan, NS & Goodson, JL (1978) The Plan Recognition Problem: An intersection of psychology and artificial intelligence *Artificial Intelligence, Vol 11, pp45-83*

Simon, HA (1981) *The Sciences of the Artificial* MIT Press, Cambridge, MA

Tsotsos, JK (1992) On the Relative Complexity of Active vs Passive Visual Search *International Journal of Computer Vision, Vol 7(2), pp127-142*

Wood, S (1993) *Planning and Decision-Making in Dynamic Domains* Ellis Horwood, Chichester, W.Sussex

Wood, S (1995) When Being Reactive Just Won't Do, *Proceedings of the AAAI Spring Symposium Workshop on Integrated Planning Systems*, Stanford University, March 1995