Reactive Combination of Belief Over Time Using Direct Perception

Robin R. Murphy   Dale K. Hawkins
Dept. of Mathematical and Computer Sciences
Colorado School of Mines
Golden, CO 80401

Marcel J. Schoppers
PO Box 2111
Robotics Research Harvesting
Redwood City, CA 94063

Abstract

One issue for autonomous mobile robots operating in unknown, or partially known, domains is how to handle uncertainty in their sensor observations over time. Methods such as probablistic belief networks and survivor functions are generally unsatisfactory because they require explicit models of the robot's interactions with its environment, including possible contravening events. This information is difficult to obtain, and is philosophically incompatible with reactive behaviors.

This paper presents an approach which eliminates the need for explicit models and reasoning; instead, it relies solely on directly perceivable attributes of the robot, object, and environment. The attributes qualitatively rate whether the robot's current observations are from an inherently more informed state than previous readings (e.g., from a better viewpoint). Observations from more informed states have different rates for the accrual and attrition of belief than those taken from less informed states. This paper describes the implementation, focusing on how the information state is computed using fuzzy logic, and how the state dynamically adapts a variation of Dempster's rule to generate the total belief. Data from a mobile robot tracking an unknown object demonstrates that the reactive computation of belief over time performs well for six canonical accrual and attrition cases.

1 Introduction

One issue for autonomous mobile robots, and other agents, is how to handle uncertainty in its observations about an object or scene over time. Combining uncertainty over time is problematic for formal theories of evidential reasoning, i.e. Bayesian and Dempster-Shafer, because the current observation is not independent of the previous one. In order to generate a reasonable total belief, some means of combining the past belief with the instantaneous belief must be applied. However, such a method must take into account contradictory influences; for example, older belief tends to be less believable [Drainkov and Lang, 1993], but objects tend to persist [Dean and Wellman, 1991].

Methods such as dynamic belief networks [Dagum *et al.*, 1992; Dean and Kanazawa, 1989] and survivor functions [Dean and Wellman, 1991] are unsatisfactory for autonomous mobile robots operating in unknown domains, both in theory and in practice. These methods assume explicit *a priori* models of the object's behavior, interactions with the environment, and possible contravening events in order to update the relationship between past and present belief. A robot operating in a partially known environment, e.g. Mars, or in unpredictable situations, e.g., construction of the space station, may not have access to correct models. If the robot perceives something unusual, it will need to track or continuously observe it in order to determine if it is a hallucination without necessarily knowing what it is. Also, the adaptation of a probablistic belief network to a new situation involves explicit reasoning on a global level. This is incompatible with mobile robots which use a reactive or a hybrid deliberative/reactive architecture. In these architectures, sensing at a reactive level involves only local behavior-specific representations (if any) and rapid update times. Attempting to integrate global models and reasoning into the behaviors both violates the principles of the architectures and could slow down execution. Therefore, a mechanism is needed which can reactively adapt the combination of belief.

This paper presents an approach to this problem motivated by the work on direct perception done by the cognitive psychologist J. J. Gibson [Gibson, 1979]. Gibson maintains that there are naturally occurring *affordances* for each activity of an agent. An affordance is a perceivable potentiality of the environment that supports the intended action without requiring memory, inference, or interpretation. One well-known example of an affordance is the use of optic flows to compute the time-to-contact with an object without having to identify the object.

By using direct perception, the combination of belief over time eliminates dependence on explicit models about the object or environment or global reasoning about events. Furthermore, it is independent of the recognition process. The application of the principles of direct perception to the combination of belief is conceptually similar to the shift in robotics from deliberative mechanisms to reactivity, most notably, reactive behaviors and reactive planning.

This paper describes the implementation of the direct perception approach using a form of Dempster-Shafer theory, and fuzzy logic to compute the agent's information state. The paper presents data from a mobile robot showing that the reactive computation of belief over time follows intuition for six canonical belief cases captured by three scenarios.

2 Approach

Our approach has two central tenets. First, the *information state* of the agent relative to the object should determine the combination of belief about the object over time. For example, if the agent is moving towards the object and is accurately tracking it, then its current observation is now *inherently more informed* than its previous observations; the belief in the object from the current information state should count more than previous, less informed observations. Note that the term "belief in the object" is used to mean belief that the object is what the robot thinks it is; i.e., that it is the same object it has seen in previous observations. Likewise, if the agent is moving away from the object or cannot stay focused on the object, then it has an inherently less informed vantage and so the past should weigh more than the present.

The information state should impact the rate of decay in belief in an object. If the object disappears, and the sensors are indeed working, then the agent's total belief that the object was there should persist for some period of time, allowing it to continue its behavior. There are many situations where the object might temporarily disappear but are difficult to reason about at the reactive level: the object may be occluded, the agent may turn to avoid an obstacle. In many cases, the agent will be able to reacquire the object if it continues with its behavior. Therefore, it is of practical importance to have a reasonable decay rate. Our approach assumes that if an object disappeared while the agent was at an informed state, the belief should persist for longer than if it was at an uninformed state. If the agent wasn't sure of what it was seeing (low initial belief), and did not have a particularly good vantage point, then the belief would start low, decay rapidly, and the agent would declare the object missing more quickly than if it started with a high certainty in the object from a good vantage point.

The second tenet states that the information state can be directly observed, independently of the belief about the identity of the object of interest. For example, an extended Kalman Filter (EKF) can be used to track an unknown object as long as it has some distinguishable features, e.g., color. The extended Kalman Filter has some measure of uncertainty, its tracking error; this uncertainty reflects how well the agent is inherently able to perceive the object and make a correct identification (or perform the desired task). Uncertainty about the quality of sensing determines the information state of the agent relative to the object, while uncertainty in the object forms the belief which is combined over time based on the change in information states.

3 Implementation

[Murphy, 1996] has demonstrated an adaptive rule of combination which adapts the belief updating process based on a *contextual weighting parameter,* n, which can be used to represent the information state of the agent-relative to the object. However, that work did not fully develop a mechanism to compute n based solely on direct perception, nor did it demonstrate the impact of the rule for representative characteristic scenarios of belief accrual and attrition. In this paper, n represents the information state of the agent relative to the object. This section summarizes the adaptive rule of combination, and describes the computation of n using fuzzy logic and directly perceivable attributes of an unknown object. The reader is directed to [Murphy, 1996] for a more detailed discussion of Dempster-Shafer theory and the adaptive rule of combination.

3.1 Adaptive Rule of Combination

[Smets, 1991; Wilson, 1993; Zadeh, 1986] have argued that Dempster's rule is only one possible rule for combining Shafter belief functions, representing a set of specific assumptions. [Murphy, 1996] proposes an alternative rule of combination, capturing different assumptions: observations are not independent, missing belief should not be an identity function, and the order of combination is important. The alternative rule is:

$$m(C_k) = \frac{\sum_{A_i \cap B_j = C_k; C_k \neq \emptyset} f(m(A_i) \times m(B_j))}{\sum_{A_i \cap B_j \neq \emptyset} f(m(A_i) \times m(B_j))} \qquad (1)$$

where:

$$f(m(A_i)m(B_j)) = [m(A_i) \times m(B_j)]^n, 0.0 < n < 1.0$$

Note that when n = 1.0, the rule degenerates to Dempster's rule of combination.

This rule has been empirically shown in [Murphy, 1996] to change the rate of accrual and decay of total belief from old observations compared to current belief (new observations) according to the value of n. The output of this rule differs from Dempster's rule in two ways. First, the old observations are no longer treated equally with the current observation. The value of n determines whether the past or the current observation dominates the combination of belief. But the rule does not just change the relative weighting of the past with present, as the decay function in [Guan and Bell, 1993] would. Instead, it modifies the rate of accrual of belief for the object and the rate of attrition separately. Lower values of n cause the total belief in the landmark to accrue supporting observations slowly (past belief dominates or smooths out the accrual rate) but decay quickly

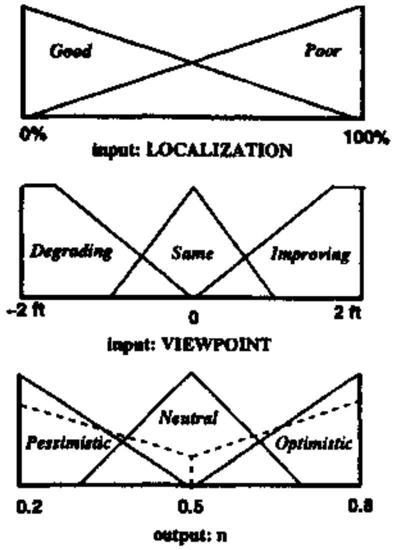

Figure 1: Fuzzy set partitionings/rnembership functions for LOCALIZATION, VIEWPOINT, and *n*. (Dashed lines show impact of GENERALLY hedge).

in the presence of missing or revising beliefs (current belief dominates decay rate). Values approaching 0.0 produce *pessimistic* combinations of belief (tends to believe the worst). Values approaching 1.0 are sensitive only to supporting observations, accruing belief quickly and decaying slowly; this is to be expected since this is converging on Dempster's rule. These high values of *n* will be referred to as *optimistic.* Values around 0.5 produce a conservative or *neutral* behavior, with slow accrual and slow decay.

When the combination of belief produces a decay in belief for a proposition, the discounted belief mass is distributed to the other propositions. This produces an averaging effect. If the agent is capable of observing only two propositions (as in the demonstrations), an unchecked decay in belief will lead to all propositions having a value of 0.5. This condition is equivalent to the "vacuous" belief function in Dempster-Shafer theory, where the belief function contains no information.

3.2 Computation of n

At this time, the information state n is computed as a function of two directly perceivable influences: *LOCALIZATION* and *VIEWPOINT.* LOCALIZATION captures the accuracy of the agent's ability to track and localize the object using the EKF. VIEWPOINT indicates whether the agent is in an inherently more informed viewpoint to observe the object.

Fuzzy logic was chosen to represent and combine influences for several reasons. Belief from real-time sensor observations is well-suited to a fuzzy representation, since the absolute value of the belief at any given time is not critical, just the trend. Also, fuzzy logic permits the addition of other variables and easy modification of the rules. Another potential advantage not explored in this paper is the use of fuzzy set hedges (very, somewhat) to modify the output based on *a priori* knowledge or motivation. For example, if the agent Ls highly motivated to find an object, the overall process should be very optimistic, resulting in a slow decay rate or a prolonged persistence of belief.

The fuzzy sets for LOCALIZATION, VIEWPOINT, and n are shown in Fig. 1. LOCALIZATION is partitioned into GOOD and POOR. The domain for LOCALIZATION is the percent error in the EKF compared to the field of view of the sensor. If the agent is unable to track the object, the region where the object could be will grow to exceed the field of view of the sensor, e.g., 100% error. If the expected occurrence of the object coincides with the actual observation of the object, the agents will be perfectly tracking the object leading to 0% error. In the demonstrations, computer vision was used to track an unknown object based on a color histogram. The EKF computed the expected locations of the upper left and lower right corners of a bounding box.

The VIEWPOINT variable consists of three sets: DEGRADING, SAME, and IMPROVING. The measurement of the inherent quality of the VIEWPOINT depends on the sensor(s) and feature extraction algorithms(s), For the sake of simplicity, it was computed in the demonstrations, as the change in distance (in feet) between the robot and the object since the last observation based on shaft encoder data. Other methods (e.g., optical flow, stereo range) could have been used to estimate the change in quality of the viewpoint.

The output of *n* was divided into three fuzzy sets: PESSIMISTIC, NEUTRAL, and OPTIMISTIC, *n* Ls computed from rules encapsulated in the Fuzzy Associated Memory shown below:

Localization	*Viewpoint*		
	degrading	*same*	*improving*
good	optimistic	generally optimistic	pessimistic
poor	neutral	generally pessimistic	neutral

The fuzzy hedge GENERALLY was applied to PESSIMISTIC and OPTIMISTIC sets. GENERALLY diffuses the set, resulting in a more neutral output value of *n.* The rules attempt to express that if the agent is moving (viewpoint is not the same) and is unable to track the object, it should react conservatively to any new observations because the instantaneous belief may be the result of temporary occlusions which can disrupt tracking. However, if the agent is stationary and still cannot track the object, low instantaneous belief could be a sign that the wrong object has been acquired or higher quality sensors and/or algorithms should be employed. Therefore, poor localization from the same viewpoint makes *n* generally pessimistic. If tracking is good, and the agent is moving away from the object, the instantaneous belief will generally get worse due to the degrading viewpoint, n is set to be optimistic in order

Figure 2: Denning MRV4 robot tracking the object (Tweety and Sylvester).

to smooth out the change in total belief by dampening the decay rate. Likewise, if the agent is moving towards the object and tracking well, n is set to pessimistic in order to react to any instantaneous belief that the object is not the object that the agent was looking for. In the case where the robot is tracking the object but is stationary, the generally optimistic value of *n* allows the total belief to accrue somewhat due to repeated observations. Defuzzification of the fuzzy rules is performed by the computing the centroid (composite moments) of the output sets.

4 Demonstrations

In this section, the use of direct perception for combining belief over time is demonstrated with a mobile robot tracking an object for six canonical cases. The six cases represent the spectrum of belief updating and revision activities following [Dean and Wellman, 1991]): *natural accretion, natural attrition, causal accretion, causal revision, spontaneous causation,* and *spontaneous revision.* The six cases are arranged into three scenarios, pairing the cases.

The demonstrations were performed on a Denning MRV4 research mobile robot using a reactive *move- to - goal* behavior. The behavior used an onboard camcorder and framegrabber to track the goal, a cardboard poster of Tweety and Sylvester shown in Fig. 2. The robot was presented with the poster; it constructed a color histogram model [Swain and Ballard, 1990] from an initial observation, then began to track it. The Shafer belief function representing the belief for the object was computed as the percent intersection of the model histogram with the observed histogram *(m(P)).* The object image size was scaled to match the model. The percent empty intersection became uncommitted belief. This was done because the robot could not discern whether the mismatch was due to occlusion, which meant the belief mass associated with mismatch should be plausible, or was due to seeing features of another object, n was computed using the fuzzy sets described previously. Because the demonstrations rely on a real robot operating in real-time, the data may appear to be cluttered and/or misleading. The reader is directed to [Murphy, 1996] where the behavior of the adaptive rule for fixed values of *n* is described. Also, the total belief using Dempster's rule is not shown on the data sets to reduce clutter. In each scenario, natural accretion drives the total belief in the object to complete belief (1.0) in short order.

It should be emphasized that the *move-to-goal* behavior does not rely on *a priori* models. The robot can be initially presented with any object with a color signature distinguishable from the walls. The demonstrations have been repeated with different objects (signs, posters) and reacted with the same trends in the combination of belief. Also, the color histogram could be substituted with a more rigorous recognition algorithm; this would change the belief in the object, but not affect the direct perception of the information state and reactive combination of belief over time.

4.1 Natural Accretion and Attrition

Natural accretion is the tendency for belief in an object to accrue when there is no other source of belief or contravening event. For example, if the observer is stationary, the object is stationary, and the belief in the object at each observation is the same, then a limited increase in the total belief is reasonable due to its repeated observation. The counterpart to accretion is attrition. If there is no source for belief about an object (e.g., missing observations), the total belief should decay.

In this demonstration, both the robot and object were stationary. The object was initially present (accretion phase) for 18 updates, then removed to generate missing observations (attrition phase) for 5 updates, then reintroduced and removed at different intervals. Finally, the poster is removed for an extended period of time (after 70 updates). Fig. 3 shows the current belief, the total belief, and the value of *n* computed by the fuzzy controller from the tracking error and change (delta) in position. The variation in the tracking error is from sensor noise.

The graph shows that the total belief does accrue somewhat, then levels out. It also shows that the total belief decays over time. In the final removal, the tracking error grew to 100%, and the belief did decay to 0.5, the equivalent of the vacuous belief function for the alternative rule of combination [Murphy, 1996]. If the belief had been combined with Dempster's rule, it would have almost immediately accrued to 1.0 (complete belief) for the object, and any attempt to revise it would have been ignored by the rule of combination.

4.2 Causal Attrition and Accretion

Causal accretion and attrition is the change in belief due to some external contravening event, either favorable or otherwise. This scenario duplicated "tunnel" experiments from child psychology, where the object moved behind a screen and reappeared on the other side. In the

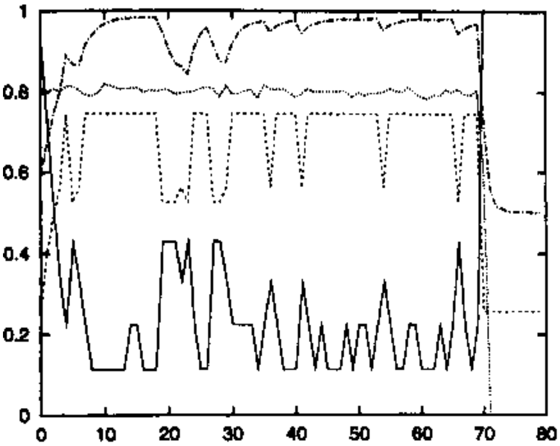

Figure 3: Behavior of belief for natural accretion and attrition scenario. Tracking error: solid line, delta position: long dashed line, N: short dashed line, current belief: dotted line, total belief: dash-dot line.

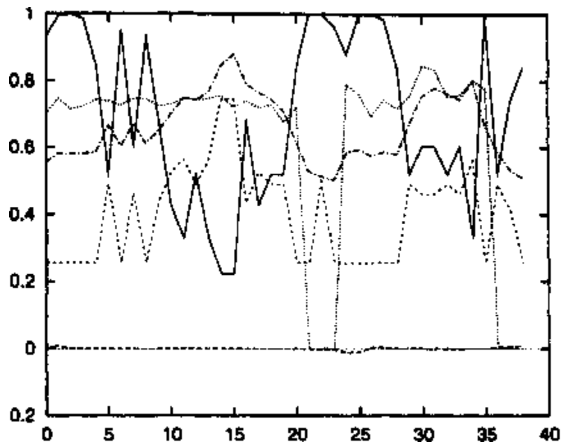

Figure 4: Behavior of belief for causal attrition and accretion scenario, robot stationary, object moving.

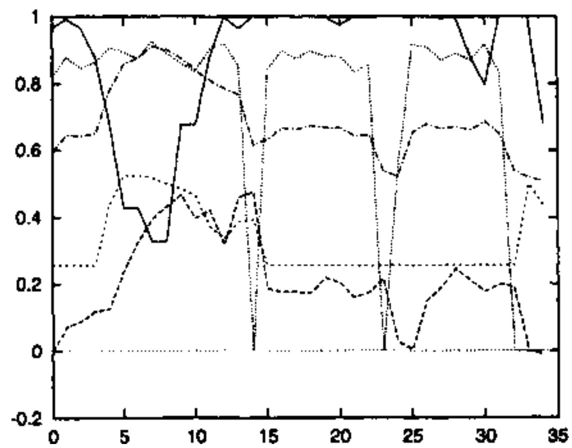

Figure 5: Behavior of belief for causal attrition and accretion scenario, both robot and object moving.

first demonstration, the robot was stationary to reduce the impact of tracking errors due to the combined movements of the object and robot. In the second demonstration, the robot was allowed to move to the object, essentially playing tag with it. Previous work [Murphy, 1996] has shown that different values of *n* (not computed dynamically) will change the decay rate, thereby changing whether the robot will abandon the tracking before the object emerges on the other side of the screen. The purpose of these two demonstration was to show the reactive computation of n and the subsequent impact on robot behavior. It should be noted that by the use of direct perception, the robot did not have to identify that a contravening event had occurred in order to combine belief in a reasonable manner.

The results of the first demonstration are shown in Fig. 4. The data shows the object moving but being visible for 20 cycles, then disappearing behind the screen for 7 updates, reappearing and then disappearing behing the screen permanently. Although the noise in the tracking error makes this graph more difficult to read, it can be seen that the total belief decays significantly when the poster disappears, but not so much that it gives up tracking. If the object remains hidden, as at the end of the scenario, the total belief will reach 0.5 or vacuous, and the robot will terminate the behavior.

The results of the second demonstration are shown in Fig. 5. The robot was presented with a new, smaller object (a green poster board target) to allow it to get closer to the target before stopping due to looming. The robot is able to successfully track the object each time it reappears on the other side of the screen, because the total belief persists.

4.3 Spontaneous Causation and Revision

Spontaneous causation and revision occur when the belief improves or degrades without a contravening event, in effect, due to the normal progress of the activity. In the demonstration shown in Fig. 6, the robot moved toward a sign with a color histogram similar to the Tweety and Sylvester poster. As it moves toward it, the current belief should be going up since the the histograms should have a better match. This improvement in viewpoint is shown on the graph, and as per the fuzzy rules, drives the value of n pessimistic (low). Therefore, the lack of improvement in current observations actually serves to decrease the total belief in the object being Tweety and Sylvester to the vacuous state, and the robot correctly terminates the *move-to-goal* behavior and turns to search for the poster. Note that the combination of belief allowed the robot to determine that it was being fooled even though at each instant, the current belief was fairly

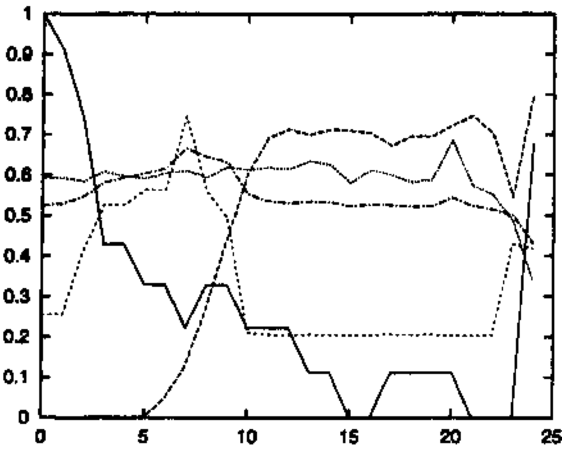

Figure 6: Behavior of belief for spontaneous attrition scenario.

good.

The case of spontaneous accretion, where the robot is able to maintain high belief as it recedes from an object, is not shown here, because the noise in the tracking errors make the graph hard to read.

5 Summary and Conclusions

This paper has shown how two affordances (e.g., localization error and viewpoint) can be used to estimate the change in the information state of belief of object observed by an autonomous mobile robot. This change in information state modifies the rate of belief accretion and attrition, as calculated with a variant of Dempster's rule of combination. This method eliminates the need for explicit, *a priori* models of the object's projected behavior over time, the sensitivity of the robot's sensors, and the impact of the environment on perception. Another advantage of the direct perception approach is that it is reactive. As demonstrated with a mobile robot, it can be used in real-time with a behavior relying only on local representation of the object (i.e., a color histogram).

The use of affordances allows belief updating and revision for objects which have never been seen before, and so have unknown properties. Furthermore, since the method uses direct perception, it does not require any reasoning about possible contravening events introducing sudden changes in belief (e.g., occlusion, sensor failure, etc.). Instead, the belief combination process incorporates the impact of the belief (e.g., low or missing belief). If the contravening event persists, the total belief will either improve or degrade over time. This total belief is used by the robot to determine what to do next; if necessary, the robot can switch from a reactive control mode to a deliberative regime.

Current work is concentrating on an integrated implementation with the NASA/JSC robot simulator, providing more behaviors, more sensors, and scenarios. We are also examining the use of fuzzy set hedges to represent other external factors influencing the combination of belief, such as the motivation of the robot. Although this method has been demonstrated only for the domain of an autonomous mobile robot operating in unknown, or partially known, environments, we believe it can be extended to other instances of situated agents.

Acknowledgments

This work is supported in part by NASA/JSC Contract NAS.9-19040, NSF Grant IRI-9320318, and ARPA Grant AO#B460. The authors would like to thank Ken Hughes, Daniel Shapiro, Glenn Blauvelt, Dave Hershberger, and the anonymous reviewers for their helpful comments.

References

[Dagum *et al,* 1992] P. Dagum, A. Glaper, E. Horvitz. Dynamic network models for forecasting. In *Eighth Conference on Uncertainty in Artificial Intelligence,* pages 41-48, 1992.

[Dean and Kanazawa, 1989] T. Dean and K. Kanazawa. A model for reasoning about persistence and causation. *Computational Intelligence,* 5(3): 142-150, 1989.

[Dean and Wellman, 1991] T.L. Dean and M.P. Wellman. *Planning and Control.* Morgan Kaufmann, San Mateo CA, 1991.

[Drainkov and Lang, 1993] D. Drainkov and J. Lang. Possibilistic decreasing persistence. In *Ninth Conference on Uncertainty in Artificial Intelligence,* pages 469-476, 1993.

[Gibson, 1979] J.J. Gibson. *The Ecological Approach to Visual Perception.* Houghton Mifflin, Boston MA, 1979.

[Guan and Bell, 1993] J.W. Guan and D.A. Bell. Discounting and combination operations in evidential reasoning. In *Ninth Conference on Uncertainty in Artificial Intelligence,* pages 477-484, 1993.

[Murphy, 1996] R. Murphy. Adaptive rule of combination for observations over time,. In *1996 Multisensor Fusion and Integration for Intelligent Systems (MFI96),* pages 125-131, 1996.

[Smets, 1991] P. Smets. About updating. In *Seventh Conference on Uncertainty in Artificial Intelligence,* pages 378-385, 1991.

[Swain and Ballard, 1990] M. Swain and D.H. Ballard. Indexing via color histograms. In *International Robots and Vision Automation Conference,* pages 378-385, 1990.

[Wilson, 1993] N. Wilson. The assumptions behind dempster's rule. In *Ninth Conference on Uncertainty in Artificial Intelligence,* pages 527-534, 1993.

[Zadeh, 1986] L.A. Zadeh. A simple view of the dempster-shafer theory of evidence and its implication for the rule of combination. *AI Magazine,* Summer:85-90, 1986.

ROBOTICS

Robotics 2