

.....

Received 23 January 2010; accepted 16 August 2010

Journal of Field Robotics 27(6), 819–833 (2010) © 2010 Wiley Periodicals, Inc.
View this article online at wileyonlinelibrary.com • DOI: 10.1002/rob.20371

Report Documentation Page				Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE JAN 2010		2. REPORT TYPE		3. DATES COVERED 00-00-2010 to 00-00-2010	
4. TITLE AND SUBTITLE 360-Degree Visual Detection and Target Tracking on an Autonomous Surface Vehicle				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) California Institute of Technology, Jet Propulsion Laboratory, Pasadena, CA, 91109				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Same as Report (SAR)	18. NUMBER OF PAGES 15	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

vicinity or (b) a previously detected target has disappeared from its previously confirmed position. Information regarding the alert target is passed to an autonomous inspection vehicle and also to the human command and control team.

- For the high-value asset protection (HVAP) mission, the ASV patrols a littoral or offshore region, circling a fixed asset. SAVAnT monitors the surrounding sea for targets that may be approaching the asset. Upon identifying a target with sufficient confidence and within a given range, the ASV deviates from its path to approach the target boat for closer inspection. Here, the patrol region is much smaller than in the AFP mission but there is greater uncertainty about the direction from which the target may be approaching.

Autonomous perception and planning for these maritime applications is difficult for numerous reasons. Most challenging is the task of real-time visual detection of particular three-dimensional (3D) objects (primarily certain boat types) in an image under widely varying conditions, affecting viewing angle, lighting, partial occlusions, range, background, etc. Downstream of the detection process, the tracker attempts to estimate the objects' states based on sequences of noisy, bearings-only measurements and must account for false positives, missed detections, and possible occlusions. Finally, the autonomous planning and navigation system must include algorithms for reasoning through the appropriate reactions to the surrounding contacts, using the (imperfect) output of the perception system to maximize chances for mission success.

The AFP mission's tracking problem is especially difficult, as long periods of time might elapse between viewing the same target and SAVAnT must maintain the target's location and identity over time to confirm that the target still exists or to declare its absence. Note that there may be many changes in the visual scene that are *not* of interest on successive passes—for example, different “friendly” boats (clutter) may be nearby, the lighting conditions and weather may have changed, and the ASV's perspective of the fixed asset may be different. Further, the target may be “hiding” among much larger boats or along a populated shoreline, making it viewable only from certain angles for a short time. Despite these challenges, we wish to have a reliable alert system to identify changes in our targets' existence or positions. To our knowledge, this is a problem that has not been previously addressed.

The realities of the maritime environment further complicate our scenarios. The pitching and rolling of the ASV affects the perceived motion of objects. Reflections off the water surface can cause confusion, and at low sun angles surface objects are often recognizable only by silhouette. Although wakes can be useful for identifying boats in a scene, they are of course absent when the target is stationary, which plays a role in deciding how to train detection algorithms. Finally, image differencing is not readily used

as a tool because it would result in many false positives on surface waves, and medium-to-high sea states can result in occlusions and false positives on white caps.

To address these challenges, the SAVAnT system decomposes the contact detection and target tracking tasks. We first apply contact detection algorithms, as well as several helpful image preprocessing steps, separately to each image of the six cameras that provide the 360-deg panoramic view, providing the bearings of detected contacts in a frame. Then, a data fusion and target tracking module analyzes these measurements, hypothesizing the correct data associations, rejecting false positives, and estimating target position. Alert conditions are identified via a novel method that estimates each target's probability of “existence.”

Work reported by Benjamin, Curcio, Leonard, and Newman (2006) gave details of successful in-water demonstrations of a behavior-based system that had the rules of the road explicitly built into the behavior base. In their demonstrations, position information was shared directly between vehicles via a wireless link, rather than needing to be perceived by individual ASVs. Larson, Ebken, and Bruch (2006) recently reported on a behavior-based hazard avoidance (HA) system for ASVs that combines deliberative path planning with reactive response to close-in dynamic obstacles. Their system used digital nautical charts (DNCs) for the initial long-range path planning coupled with a passive stereo system developed at JPL for the reactive control. The current implementation of their system does not link any resource use-based planning into the HA behavior.

An active sensor approach to hazard detection using a laser range finder for navigation was reported by Jimenez, Ceres, and Seco (2004). Snyder, Morris, Haley, Collins, and Okerholm (2004) successfully demonstrated the components needed for autonomous in-water navigation in harbor and riverine environments. Their system used six cameras arranged as a 360-deg color sensor coupled with sky/sea/shoreline segmentation, optic flow, and structural model techniques to determine the relative position of obstacles and safe paths. Nervegna and Ricard (2006) described their simulation work in higher level command and retasking of multiple heterogeneous air, surface, and underwater vehicles. Their risk-aware, mixed-initiative dynamic replanning (RMDR) system uses a mixed initiative interaction module (MIIM) for the operator interface and a dynamic replanning and situation assessment (DRASA) for onboard autonomous control.

Most of the works cited above have concentrated on the navigation and HA aspects for control of autonomous boats. More complicated scenarios such as HVAP are based on mission-level autonomy that incorporates these baseline aspects into an integrated approach. The RMDR system addresses the planning aspects of such a mission, but the onboard autonomy is assumed to already be in place. CARACaS explicitly includes the blend of sensing and

behavior-based autonomy to build the more complicated mission scenarios.

Additionally, the patrol ASV scenarios we address here present new challenges, particularly for the AFP mission. In most traditional multitarget tracking scenarios, a fixed sensor continually monitors a given sensor volume; our region of interest is much larger than our sensor range, and so the problem relies on the movement of the sensor through this region for coverage. Additionally, we may leave the sensor range of a target and return to it later, wishing to estimate whether it is still the same target.

Finally, the AFP mission might also be viewed from the perspective of visual change detection, comparing the scene around the fixed asset on successive passes through the patrol region. For example, Perera and Hoogs (2004) offer a change detection solution that operates on an “object level,” as ours does. However, we note that several aspects of our problem differ from those addressed by these and other authors. First, we wish to detect only certain changes, ignoring motion of “benign” vehicles and shore activity (and of course clouds and waves). Second, we must allow for significantly different camera positions on separate passes. Also, we wish to track targets to estimate their location. Finally, we note that our method requires neither image registration nor training data of empty scenes.

The remainder of this paper is organized as follows. We present an overview of autonomy system, including its capabilities and architecture, in Section 2. We detail the core components of SAVAnT—including contact detection and target tracking/change detection—in Section 3. We describe our on-water experimental setup and test scenarios in Section 4 and the corresponding results in Section 5. Finally, concluding remarks are given in Section 6.

2. AUTONOMY SYSTEM ARCHITECTURE

2.1. CARACaS Overview

Several key aspects of an intelligent autonomy approach to ASV control include the handling of the inherently uncertain nature of dynamic sea surface operations; sensing for hazard detection/avoidance and situational awareness; behaviors for obeying the rules of the road during interactions with other manned and unmanned vehicles; cooperation among heterogeneous vehicles on the sea surface as well as underwater and in the air; onboard resource-based planning for mission operations; integrated system health maintenance for long-duration missions; and the human operator command interface. JPL has developed a tightly integrated instantiation of an autonomous agent called CARACaS, a block diagram of which is shown in Figure 1, to address many of the issues for survivable, autonomous ASV control (Hansen, Huntsberger, & Elkins, 2006). CARACaS is composed of a dynamic planning engine, a behavior engine, and a perception engine. The




Figure 1. Block diagram of CARACaS. The network in the behavior engine is built from primitive (dark gray) and composite (light gray) behaviors. The dynamic planning engine interacts with the network at both the primitive and composite behavior levels.

SAVAnT system is part of the perception engine, which also includes a stereo vision system for navigation.

2.2. Dynamic Planning Engine

The dynamic planning engine leverages the CASPER (Continuous Activity Scheduling Planning Execution and Re-planning) continuous planner (Chien, Knight, Stechert, Sherwood, & Rabideau, 2000; Chien, Sherwood, Tran, Castano, Cichy, et al., 2003) developed at JPL. Given an input set of mission goals and the autonomous vehicle’s current state, CASPER generates a plan of activities that satisfies as many goals as possible while still obeying relevant resource constraints and operation rules. [CASPER has been used to autonomously perform the planning/replanning for the Earth Observation 1 (EO1) satellite continuously since November 2004.] Plans are dynamically updated using an iterative repair algorithm that classifies plan conflicts (such as a resource over-subscription) and resolves them individually by performing one or more plan modifications.

2.3. Behavior Engine

CARACaS leverages the results of previous efforts at JPL in the multiagent control architecture CAMPOUT (Control Architecture for Multi-robot Planetary Outposts) (Huntsberger, Cheng, Baumgartner, Robinson, & Schenker, 2003; Huntsberger, Trebi-Ollennu, Aghazarian, Schenker, Pirjanian, et al., 2004; Huntsberger, Trebi-Ollennu, Nayar, Aghazarian, Ganino, et al., 2003) in order to develop behavior composition and coordination mechanisms. CARACaS uses finite state machines for composition of the behavior network for any given mission scenarios. These finite state machines give it the capability of producing formally correct behavior kernels that guarantee predictable performance.

For the behavior coordination mechanism, CARACaS uses a method based on multiobjective decision theory (MODT) that combines recommendations from multiple behaviors to form a set of control actions that represents their consensus. CARACaS uses the MODT framework (Pirjanian, 2000) coupled with the interval criterion weights method (Benjamin, 2002a, 2002b) to systematically narrow the set of possible solutions (the size of the space grows exponentially with the number of actions), producing an output within a time span that is orders of magnitude faster than a brute-force search of the action space.

2.3.1. Behavior Representation

CARACaS formalizes a behavior, b , as a mapping, $b : P^* \times X \rightarrow [0, 1]$, that relates a percept sequence $p \in P^*$ and an action $x \in X$ pair, (p, x) , to a preference value that reflects the action's desirability. The percept possibly includes (processed or raw) sensory input (for example, the appearance of a new contact, with related position estimate), and the N -dimensional action space is defined to be a finite set of alternative actions. The described mapping assigns to each action $x \in X$ a continuous valued preference, where the most desired actions are assigned 1 and undesired actions are assigned 0 from that behavior's point of view. In CARACaS, behaviors are activated using simple, two-state finite state machines ("idle" and "run"), in order to maintain real time control over which collection of behaviors are active at any given time in a deterministic way.

2.3.2. Behavior Composition

Behavior composition refers to the mechanisms used for building higher level behaviors by combining lower level ones. A major issue in the design of behavior-based control systems is the formulation of effective mechanisms for coordination of the behaviors' activities into strategies for rational and coherent behavior. Behavior coordination mechanisms (BCMs) manage the activities of lower level behaviors within the context of a high-level behavior's task and objectives. For a detailed overview, discussion, and comparison of behavior coordination mechanisms, see Pirjanian (1999).

CARACaS predominantly uses the primary sequential and parallel composition operators, represented as $\Rightarrow$ and $\parallel$, respectively. A simple example of the use of the sequential composition operator is that used in the HVAP mission behavior:

Patrol $\Rightarrow$ Intercept $\Rightarrow$ Inspect $\Rightarrow$ Patrol . . . ,

where each individual high-level behavior completes before the next one starts. A simple example of the use of the parallel composition operator is that used in the go-to-waypoint with HA behavior:

Maintain_Track $\parallel$ Avoid_Hazards,

where the parallel composition operator $\parallel$ can be any number of BCMS that are used to coordinate the activities. Among the two used most often are the AND and OR operations that can be defined in a number of different manners.

The simplest definition for the OR composition operator is mutual exclusion, meaning that either Maintain_Track is generating the rudder and throttle commands or Avoid_Hazards more or less independently. In the case of the AND composition operator, both of the behaviors are contributing to the commands for the generation of action. Among the most common blending or fusion methods for the use of an AND operation for fusion are voting techniques (Huntsberger & Rose, 1998), fuzzy (Saffiotti, Konolige, & Ruspini, 1995; Yen & Pfluger, 1995), and MODT (Pirjanian, 2000). CARACaS currently builds the fusion of the two behavior outputs into the finite state machine used for the Avoid_Hazards behavior by biasing the choice of safe paths for navigation toward the waypoint goal, thus accommodating the Maintain_Track behavior at the same time. For details of the underlying behavior-based framework for CARACaS, see Huntsberger, Trebi-Ollennu, et al. (2003).

3. SAVANT PERCEPTION ENGINE

The components of the SAVANT system are depicted in Figure 2. SAVANT receives sensory input from an inertial navigation system (INS) and six cameras, which are mounted in weather-resistant casing (see Figure 3), each pointed 60 deg apart to provide 360-deg capability, with 5-deg overlap between each adjacent camera pair. The core components of the system software are as follows. The *image server* captures raw camera images and INS pose data and "stabilizes" the images (for horizontal, image-centered horizons). The *contact server* detects objects of interest (contacts) in the stabilized images and calculates absolute bearing for each contact. The *OTCD server* (object-level tracking and change detection) interprets series of contact bearings as originating from true targets or false positives, localizes target position (latitude/longitude) by implicit triangulation, maintains a database of hypothesized true targets, and sends downstream alerts when a new target appears or a known target disappears (see Section 3.2).

3.1. Image and Contact Servers

As noted above, the combined goal of the image server and contact server is to process the camera images to detect objects of interest (i.e., contacts), resulting in the bearing measurements of the contacts in each frame.¹ Searching the images for objects of a particular type is the biggest challenge of these modules—the already-difficult task of real-time

¹In this paper, the term *frame* is used to describe a set of (six) images taken from all cameras at the same time.




Figure 2. System components and data flow of the SAVAnT system.

object recognition is complicated by highly variable lighting conditions, arbitrary viewing angles, possible occlusions, different contact ranges, and possibly high sea state. Such variability is illustrated by the example images of desired contacts in Figure 4. The top left image shows the target boat sitting in front of a large ship that represents the asset to be protected. This was recorded on a sunny afternoon with wide dynamic range in the scene, so the camera gain was low and the target boat has low contrast with the ship behind it. Also, the camera is oriented at an angle away from horizontal, due to its mounting, the ASV riding angle (which depends on its speed), and the current pitch and roll of the ASV. The top right image shows the same scene from the same camera on a cloudy day. Here the camera gain is high and the target boat has greater contrast with the ship behind it. In the bottom left, the scene is shot from another angle late in the afternoon with the sun setting behind the ship. This image presents two difficult conditions: the target boat lies in the shadow of the ship, and the sun

glare reflects from the water and waves directly into this camera view. All three of these images also contain clutter from the shoreline behind the ship, including small boats and buildings that approximate the scale of the target boat. Finally, the bottom right image was recorded in the early evening with the camera facing away from the sun. The camera exposure time (1–2 ms) could not be increased to further brighten the image because the camera has a rolling shutter—increasing the exposure time while the ASV was pitching or rolling (or otherwise moving the camera) causes the image to warp and degrades the image stabilization in later processing.

In preparation for the object detection algorithms, the image is stabilized using current INS data, negating most of the roll and pitch of the camera frame due to sea state and boat motion. After stabilization, the image is cropped to a strip centered vertically about the estimated horizon, so that only regions of interest for surface vessels are processed, reducing computation time. The image intensity is then normalized by using local averages sampled in the sky and water. These preprocessing steps aid by making the average input to the contact detection algorithms more consistent, but significant target variability will still be present, as seen in Figure 4.

The contact detection process must be sensitive enough to pick up all the real targets from the input with a relatively low false alarm rate. We apply two custom algorithms specialized for detecting contacts of the particular vessel type(s) of interest in our scenarios. For each contact identified by one of these algorithms, a corresponding absolute bearing is backcalculated using the target image location, the camera model, the IMU-camera transform, and the IMU/global positioning system (GPS) data. The bearing and image snippet are appended to the contact list to be passed to the OTCD module and/or sent to the remote viewing interfaces.

3.2. Object-Level Tracking and Change Detection

The OTCD algorithm assimilates all the contacts identified in the contact server to generate the situation awareness required by the ASV's mission. Primarily, this responsibility takes the form of generating and maintaining a list of



Figure 3. The 360-deg camera head provides the primary data input to SAVAnT.




Figure 4. Example raw images with target, demonstrating high variability in lighting conditions, background clutter, camera roll, etc. Clockwise from top left: sunny noon, cloudy noon, evening with camera facing away from sun, and late afternoon with camera facing toward sun. Light/yellow arrow indicates asset to be protected; dark/red arrow indicates target.

targets, including confirming the existence of targets (as opposed to false-positive contacts), estimating the location of targets, and recognizing the change conditions of the AFP mission. OTCD operates at the “object level” rather than in the image domain primarily because the AFP scenario requires comparison of the region’s targets over repeated visits [between which the target(s) are not in the image] and also to obviate the need to register or stitch images from disparate cameras. OTCD accomplishes its goal by essentially tracking all targets in the patrol region over time, building a database of all confirmed targets over all visited locales, with an estimated location, covariance on the location estimate, and a probability score of whether the target currently exists at that location.

Although OTCD fundamentally solves a multitarget tracking problem, there are crucial differences between the traditional multitarget problem and our missions, particularly the AFP mission. We have a *moving* sensor that is concerned about target identity over long timescales, which must cover large gaps of time during which the target is completely out of sensor range. Usually, traditional tracking scenarios use either a fixed sensor or, if moving, one concerned about short timescales; either way, they are primarily charged with monitoring *immediately visible* contacts. By contrast, we must allow the target to “leave

and come back” and still track it as the same vehicle.² Our primary innovations to cope with this challenge involve the invention of a “probability of existence” for each target as well as new ways of managing variable detection probabilities.

Note that the tracking and change detection problems would be relatively straightforward with perfect contact detection; however, because of the difficulty of the detection task, OTCD must cope with potentially heavy clutter (false positives) and with missed detections (false negatives). This is especially true because the penalty on a missed detection is high in our scenarios; thus, contact detection thresholds must remain fairly low, relying on OTCD to reject most of the false positives. In addition, OTCD must estimate two-dimensional (2D) target position (longitude/latitude) from bearings-only measurements, which are noisy due to detection of different parts of the target boat (e.g., bow vs. stern) and small imperfections in the camera models or in synchronizing the image to the IMU.

²In the current implementation, it was not practical to use identifying characteristics of an individual boat to help with this problem. However, this is a strategy slated for future work.

Finally, note that the image server and contact server operate at the image level—OTCD is the module responsible for maintaining a time history as well as for integrating the information of all six images from a single frame (e.g., managing duplicate contacts in the overlap regions and camera-to-camera handoffs).

The solution implemented in OTCD can be conceptually decomposed into three interrelated tasks. The first task is to assign incoming measurements to known targets or mark them as new targets or false positives—a classic multitarget data association exercise. The second task of OTCD is to localize the targets using multiple bearings-only measurements; an appropriate nonlinear state estimation filter can be selected that will implicitly triangulate the 2D global positions (e.g., latitude and longitude). Finally, the third task is to calculate the probability that a suspected target truly exists (in the estimated position) and use this value to determine whether alert conditions have been triggered (for new targets or disappeared targets). For this task, we have created a notion of “probability of existence,” which provides a measure that a particular target truly exists and that it is still in the location it was last observed. This probability distinguishes true targets from tracks composed of false positives (clutter) and also forms the basis of our alert conditions for new/disappeared targets.

3.2.1. Probability of Detection

The probability of detecting a target is a particularly important notion in OTCD, as it not only affects whether associations are made to a given target but also is a key factor in the probability of existence, described in Section 3.2.2. Let $P_{d,j}^k$ denote the probability that the j th existing target is detected in frame k . Unlike many traditional multitarget tracking methods, it is critical in our method that $P_{d,j}^k$ is allowed to vary per target and over time, as the contact detector’s success depends highly on the range to the target. Also, the fact that $P_{d,j}^k$ is zero when target j is out of sensor range (or occluded) allows OTCD to gracefully maintain its target list during its patrol.

To define $P_{d,j}^k$, we begin by characterizing sensor performance as a function of the range ρ , defining the nominal detection probability function $f_d(\rho)$. In our implementation, $f_d(\rho)$ is a piecewise constant function as in Figure 5. However, the true range is unknown because it is only roughly estimated by triangulation [implicitly in the extended Kalman filter (EKF)]; rather, a distribution on the range can be derived from the EKF’s estimate as a Gaussian distribution with its mean and variance ($\hat{\rho}_j^k$ and $\sigma_{\rho_j^k}^2$) determined by rotating and translating the Universal Transverse Mercator (UTM) coordinate frame to align an axis with the predicted bearing line.³ Thus, for a more accurate $P_{d,j}^k$,




Figure 5. Example functions showing the nominal probability of detection as a function of range $f_d(\rho)$ with the Gaussian distribution of a range estimate $f_N(\rho|\cdot)$ overlaid. The $f_d(\rho)$ shown in this function matches the parameters used for the HVAP mission.

we marginalize over the conditional dependence on the range:

$$P_{d,j}^k = \int_{-\infty}^{\infty} f_d(\rho) f_N(\rho|\hat{\rho}_j^k, \sigma_{\rho_j^k}^2) d\rho, \quad (1)$$

where f_N denotes the Gaussian probability density function (PDF). An example is shown in Figure 5. Note that if, in this example, we simply used the (fairly uncertain) estimated position directly, the probability of detection would be $f_d(\hat{\rho}_j^k) = 0.2$. Instead, our method results in $P_{d,j}^k = 0.45$.

3.2.2. Probability of Existence

As noted earlier, OTCD includes a new measure, which we dub the *probability of existence*, $P_{e,j}^k$, that estimates a confidence level of whether a target truly exists and is still at its estimated location. This value enables us to evaluate the AFP scenario conditions and is also used to confirm target existence in general. Let τ_j^k be an indicator variable with value 1 to signify that the hypothesized target j exists at time k and 0 indicating that it does not. Similarly, let δ_j^k indicate whether the j th target is tracked under the current data association hypothesis. After each frame, we update the probability that target j exists for every target in every hypothesis, in a Bayesian manner:

$$P_{e,j}^k = P(\tau_j^k = 1|\delta_j^k) = \frac{P(\delta_j^k|\tau_j^k = 1) P(\tau_j^k = 1)}{\sum_{\tau_j^k \in \{0,1\}} P(\delta_j^k|\tau_j^k) P(\tau_j^k)}. \quad (2)$$

$P(\delta_j^k|\tau_j^k)$ depends on the values of detection probability and of associating false positives. Let P_{FA} denote the probability that a false-positive measurement is incorrectly associated with an target track. Then the formulas we need for Eq. (2) are

that the target may be in the direction opposite the bearing), because it is just an affine coordinate transformation. In practice, this fact is inconsequential and we can define $f_d(\rho)$ for only positive ρ .

³Formally, this distribution is not on a true range; rather, this method allows negative support of the distribution (the possibility

$$\begin{aligned}
P(\delta_j^k = 1 | \tau_j^k = 1) &= P_{d,j}^k + (1 - P_{d,j}^k)P_{FA} \\
P(\delta_j^k = 1 | \tau_j^k = 0) &= P_{FA} \\
P(\delta_j^k = 0 | \tau_j^k = 1) &= (1 - P_{d,j}^k)(1 - P_{FA}) \\
P(\delta_j^k = 0 | \tau_j^k = 0) &= (1 - P_{FA})
\end{aligned}$$

The value for P_{FA} can be derived from the probability that at least one of the false-positive measurements is associated with the target; however, this value changes for every frame and for every target. We wish to make simplifying assumptions so that a constant can be used for P_{FA} without severely impacting performance. For this purpose, we make the approximation that a measurement will be associated with the j th target if it is within n standard deviations of the expected measurement. Then, the probability that a false measurements is associated with target j (i.e., one minus the probability that all N_ϕ FPs are *not* associated) is

$$P_{FA} = 1 - \left(1 - \frac{2n\sigma}{V}\right)^{N_\phi}, \quad (3)$$

where V is the observation volume and σ is the standard deviation of the innovation.⁴ To make the parameter a constant, for our application we further simplify with the approximations that $N_\phi \approx \lambda_\phi$ (the expected value used to model the probability of false positives), $\sigma \approx \sqrt{r}$ (the standard deviation of the measurement noise used in the EKF), and $n = 3$ (empirically based).

The prior $P(\tau_j^k = 1)$ is, in the simplest case, just the result from the last time step. However, this approach would not model the possibility that a target's existence can change over time (e.g., a target moves while out of range). Thus, we utilize a "forgetting factor" γ to decay the probability (up to a certain minimum threshold):

$$P(\tau_j^k) = (1 - \gamma) P(\tau_j^{k-1} | \delta_j^{k-1}). \quad (4)$$

This forgetting factor also keeps P_{FA} numerically well behaved, rather than becoming unity after many detections.

Note that there is implicit conditioning in the above probabilities that the "correct" association has been determined when tracking the target. Although there exists potential future development to link the probability of existence with the association probabilities, the above framework is sufficient for our current implementations.

4. ON-WATER EXPERIMENTAL SETUP

A series of on-water demonstrations were run at Fort Monroe, Virginia, in June 2009 (AFP mission) and in October and November 2009 (HVAP mission). A satellite

image of test zones is shown in Figure 6. The CARA-CaS/SAVANt systems were installed on two U.S. Navy experimental ASVs, pictured in Figures 7(a) and 7(b). During all of the on-water demonstrations, there was a trained operator onboard the ASV to undock and dock the boats (autonomous launch and retrieval capabilities are part of another project) and to ensure safety.

image of test zones is shown in Figure 6. The CARA-CaS/SAVANt systems were installed on two U.S. Navy experimental ASVs, pictured in Figures 7(a) and 7(b). During all of the on-water demonstrations, there was a trained operator onboard the ASV to undock and dock the boats (autonomous launch and retrieval capabilities are part of another project) and to ensure safety.

In both scenarios, a white boat, called the "PL," is used as a target boat to be identified and tracked, as shown in Figure 7(c). Although we recognize that the color of the PL is an advantage in many conditions (except when backlit), an analysis of how much of an advantage has not been undertaken. Some dark-color contacts have been trained by the same contact detectors for other applications but not in the scope of this work.

4.1. AFP Mission Setup

The AFP mission tests SAVANt's ability to recognize changes in the targets around a fixed asset during the patrol of a large region. Figure 8 shows the test layout. In our setup, the ASV travels first on a southerly pass, with the PL target docked to the hull of the fixed asset. The ASV should at this point recognize the target and send a "new target" alert. The closest distance between the ASV and the asset is between 400 and 600 m. When the ASV is more than 2 km away, the PL leaves the vicinity of the asset. Thus,




Figure 6. Satellite imagery of the field test zones.

⁴Because $\sigma \ll V$, we can be certain that $2n\sigma/V < 1$ and so this approximation is well behaved.



Figure 7. Vessels used in the field exercises.

when the ASV returns in the northerly pass it should “look for” the previously observed target and then send a “disappeared target” alert when it confirms the target’s absence. Note that in this scenario, because the ASV is concerned only with targets near the fixed asset, contacts are accepted only when the ASV is within 1,500 m of the asset and if their bearings are in a 24-deg window centered on the asset’s position.⁵ This reduces the likelihood of false positives in the crowded riverine environment; this in turn enables the detection thresholds to be set relatively low, increasing the probability of detecting the true target.

4.2. HVAP Mission Setup

The HVAP scenario tests CARACaS/SAVANt’s ability to detect a target that approaches the fixed asset in any

⁵These values are derived from the SAVANt sensor range and the size of the fixed asset plus a comfortable margin to ensure that good values are not thrown out.




Figure 8. Overhead view of the AFP mission scenario.

direction and investigate/intercept it. Figure 9 shows the scenario. The baseline case was to first respond to a drifting target just outside of the patrol zone’s perimeter, with a “stretch goal” of controlling the ASV to move on a path to intercept an (actively) incoming intruder. We present results from two tests of this scenario. In Trial 1 (22 October 2009), the PL target first drifts for several minutes about 1 km away from the asset (actually moving slowly away) and then begins approaching. Trial 2 (5 November 2009) takes place closer to the shoreline (increasing chances of false detections), and the PL drifts, staying about 800 m away from the asset.

In the HVAP mission, it is advantageous to filter out contacts that are likely to be of known objects. For example, we ignore contacts whose bearings are within 3 deg of the known fixed asset location. This technique, which eliminates the possibility of false-positive contacts arising from the fixed asset, is well justified by the scenario (we will always know the location of the fixed asset, as it even determines the patrol path, and we look for intruders approaching it). Also, we have observed several instances of false contact detections on known shoreline buildings that happen to be similar to the PL silhouette in size and outline. For the results shown in the next section, SAVANt reads in the




Figure 9. Overhead view of the HVAP mission scenario. Three boats are involved in this test: the asset boat, the ASV, and the intruder/target.




Figure 10. Probability of existence of all targets for AFP mission, with markers indicating whether the target is detected on each frame. Magnified regions are provided near the two alert events, for which the probability of detection is also shown, along with a separate plot with the estimated range to the target and the standard deviation of that estimate (ranges are on separate scales; these values have not been publicly released). The ASV is out of range of the target from about frame 85 to frame 475.

locations of selected landmarks from the region's electronic nautical chart (ENC) and removes any contact whose bearing is within about 0.6 deg of the landmark location (three landmarks were used for these trials). This provides an extensible approach to removing likely false positives with limited risk that true detections are consistently thrown out.

For the results presented in the next section, we ran SAVAnT offline in "replay mode" from the raw data logs collected during the live exercise. The distinction between live runs and replay mode is entirely transparent to SAVAnT, and the same results can be expected for a live demonstration of the same data. Illustrative results from live runs are unavailable due to technical difficulties such as communication issues and mechanical boat failures. For the results in the next section to represent an accurate assessment of the system, the system parameters have been set to be the same for all trials.⁶

⁶The AFP scenario has been constructed from two consecutive passes that actually occurred in the reverse order (northerly pass first). By switching the order, we are able to present results that test *both* types of alerts, whereas in the northerly-first order we are able to test only "new target" alert.

5. RESULTS

5.1. AFP Mission Results

In the AFP mission test (29 June 2009), the SAVAnT system correctly identified the two alert events and triggered no false positives. Viewing the targets' probability of existence (defined in Section 3.2.2) over time, as in Figure 10, provides the greatest insight into the system's hypothesized targets and how alerts were triggered. Figure 10 shows that Target 1 (the PL) exists in OTCD's target list from frames 1 to 487 and Target 2 (an unconfirmed false target) exists during frames 607–636.⁷

Let us step through the mission while examining Figure 10. Target 1's probability of existence rises quickly on the initial detections, causing the target to be *confirmed* at frame 8, generating a "new target" alert. Then, as the ASV gets closer to the target, the image size of the PL becomes too large compared to the images used to train the detectors, resulting in intervals with no detections (e.g., frames 29–45). However, because the probability of detection used by OTCD (see Section 3.2.1) is low when the ASV is close

⁷In the AFP results, each frame is separated by about 3 s.




Figure 11. Stabilized images from one frame in the AFP trial with contact detection results, showing one true positive (light/green box) and one false positive (dark/red box). The two contacts are passed on to OTCD for confirmation.

to the target, the probability of existence only slightly decreases. [The probability of detection, the range to the target (as estimated by OTCD's EKF), and the standard deviation of the range estimate are plotted in Figure 10's expanded regions.] The ASV passes the target, getting farther away and adding a few more detections (frames 50–67), but then the range to the PL begins to exceed SAVAnT's sensor range. Again, the detection probability is low at large ranges, so the probability of existence does not go to zero but instead levels out while the ASV continues its patrol away from the fixed asset.

On the return pass, the ASV comes within range of the expected target position (starting around frame 475), but the PL has moved away to an unobservable location. Because SAVAnT now expects to detect the target with non-trivial probability but receives no hits on the estimated location, the probability that it still exists at that location decreases. Finally, the probability drops below threshold, deleting the target and triggering a "disappeared target" alert.

The images in Figures 11 and 12 show example contacts of the PL target during the ASV's first pass, as well as a contact representing Target 2. Target 2 resulted from several detections of a shoreline structure or vehicle but did not trigger an alert because it stayed in *suspected* status until being shortly *deleted*. This illustrates an added layer of robustness that OTCD's status designation provides—although there were several associated hits, the target is not persistent and is correctly dismissed. Note that this is a separate mechanism from the intermittent/randomly located

false contacts that are never associated to a target but instead marked as false positives by MHT's data association process.

To detail these separate types of assignments and OTCD's performance, Table I indicates the assignments that OTCD determined for all of the contacts it received for each mission test (the HVAP results are discussed in Section 5.2). Target assignments are divided into three groups: PL (intended true target), similar boats (not the PL, but an understandable confusion), and false targets (nonboats or boats very different from the PL). The number of targets of each type is shown, with the number of contacts assigned to the targets in parentheses.⁸ "Robust misses" indicates the number of frames with no detections between the first and last detections of the true target, despite which OTCD was able to correctly maintain the target. Finally, the number of contacts OTCD marked as false positives is listed—either "correctly" (i.e., the contact indeed did not represent a target and was a spurious detection by the contact server; see examples in Figure 13) or "incorrectly" (the contact represented the PL or other target boat, but OTCD did not associate it to the target, which was rare).

⁸Recall that a *contact* is herein defined as a single detection event in one time step (frame), whereas a *target* represents a collection of these contacts that are associated over time with each other by OTCD. Thus, in the AFP scenario, SAVAnT "saw" the PL 38 times and correctly associated all 38 contacts to a single target boat.




Figure 12. Snapshots of detections in the AFP trial, with frame number. The first five images show the tracked PL target boat (Target 1). The contact in the last image was briefly a suspected targeted by OTCD (Target 2, not confirmed).




Figure 13. Snapshots of detections in the HVAP trials that were *not* the PL target boat. The boat in image (a) was a consistent detection and thus targeted by OTCD (Trial 1, Target 2), and the contacts in the other images [(b)–(c) waves, (d)–(e) shoreline buildings, (f) clouds] were correctly marked as false positives by OTCD.

Table I. OTCD assignment counts.

	AFP	HVAP 1	HVAP 2
Total contacts	66	607	440
PL targets (contacts)	1 (38)	1 (400)	1 (135)
Similar boat targets (contacts)	0 (0)	0 (0)	1 (36)
False targets (contacts)	1 (8)	0 (0)	0 (0)
Robust misses, PL	30	84	49
Correctly marked false	18	207	269
Incorrectly marked false	2	0	0

5.2. HVAP Mission Results

In both HVAP trials, SAVAnT successfully identified and tracked the target boat, as shown by the plots of OTCD’s probability of existence in Figure 14 and the target images in Figure 15. In Trial 1 (22 October 2009), the boat that OTCD labels as Target 1 is the PL (the true target), first identified at frame 38. The target is deleted at frame 233 after several consecutive frames without detections—at this point, the ASV has approached the target, causing it to be too close for the contact detection algorithms (the target is




Figure 14. Probability of existence plots for all targets identified by OTCD during HVAP mission trials. Open circles indicate frames where the contact server did not find a detection but OTCD maintained the target track. Background color indicates when target status is suspected (yellow/light gray) and confirmed (green/darker gray).




Figure 15. Enlarged snapshots of the tracked target (Target 1) from HVAP trials, with the frame number on each image. Each image is 60×60 pixels. Despite the different appearances of the targets, they are all identified by contact detectors and then tracked as the same object by OTCD.




Figure 16. Contact detection results from one frame of HVAP Trial 1. One true positive (light/green box) and two false positives (dark/red box) are superimposed on the six camera images. The false positives correspond to a lighthouse and to intermittent sun glare from a wave.




Figure 17. Contact detection results from one frame of HVAP Trial 2. One true positive (light/green box) and three false positives (dark/red box) are superimposed on the six camera images. The false positives correspond to the asset boat, a building on shore, and sun glare from a wave. The IMU-camera transform for camera 3 (bottom center) was incorrect, resulting in poor image stabilization.

“handed off” to a close-range stereo perception system). Late in Trial 1, a second target is also identified; although technically this target is a false positive because it was not the PL, it represents an accurate track of a distant sailboat whose shape matched nearly enough to the PL so that it was consistently identified by SAVAnT [see Figure 13(a)]. The ASV did not react to intercept this target, however, be-

cause it remained distant.⁹ See Figure 16 for a view through all six cameras.

⁹A heuristic system is in place to set intercept priority, based on the target’s range, persistence (number of contacts), and probability of existence.




Figure 18. ASV and target (PL) paths for the HVAP trials.

In Trial 2 (5 November 2009), SAVAnT correctly identifies and tracks the PL target, despite difficult lighting conditions (e.g., see Figure 17), for nearly the entire 500-frame trial, and the PL is still being tracked when the session is ended. Note that, because of OTCD's Bayesian-updated probability of existence, SAVAnT is robust to upstream missed detections throughout the trial, which might occur at particularly difficult temporary conditions (e.g., view angle, lighting, waves). These missed detections cause dips in the probability of existence but are not sustained enough to "delete" the target.

Images from sample false-positive contact detections are also provided in Figure 13. Note that, other than the sailboat shown in image (a), all false positives such as these are correctly marked as false contacts by OTCD and not assigned to a target. Thus, the SAVAnT output will not include these more intermittent hits. On the basis of manually scoring of the trial, we have verified there are no false associations—that is, all the detections that were assigned to the three targets (across both trials) were indeed the same boat.

The paths of the ASV (as well as the PL) for each trial are plotted in Figure 18, as collected by GPS data. In Trial 1, the ASV is nearly stationary to start and then approaches the PL as it is being tracked. In Trial 2, the ASV patrol path is more clearly visible, and the approach to the PL can be seen in the upper right. (Note that the results in Figure 14 represent only a subset of the data from the path shown, as, in our test plan, the SAVAnT system was not engaged for the entire time on water.)

6. CONCLUSIONS

This paper presented the SAVAnT autonomous perception and situation awareness elements of a patrol ASV, in the

framework of an integrated autonomy architecture (CARA-CaS). We have shown that SAVAnT can successfully detect and track other vessels in two types of asset protection missions. We have demonstrated that our contact detection methods correctly identify the target across a variety of conditions and that, by placing OTCD's additional methods of rejecting false-positive contacts downstream of the contact detector, we can set the detector thresholds lower and thus reduce the risk of missing a true target. For the challenging cases in which SAVAnT must "track" targets that are no longer in camera range for later return on the patrol, we achieved success by tracking targets in world coordinates (by decoupling the contact detection and target tracking problems) and by innovating new multitarget tracking advancements in the OTCD algorithm (particularly the handling of the probabilities of detection and of existence).

Many opportunities for future work exist. Long-range, real-time contact detection continues to be a difficult challenge, and our results included a few hypothesized targets that were not the intended boat class. Multiresolution approaches may improve these results, though there remains a fundamental question of knowing "how many pixels" are required for reliable detection of only a given target class. Narrowing the region of interest in detection—e.g., considering only image locations in which targets have been hypothesized, processing whole frames only occasionally to look for new targets—may also help by freeing computation time for further image processing. Going forward, we wish to train the contact detectors for handling multiple target classes, achieving not only detection but also classification from these algorithms.

On the tracking side, although bearings-only measurements can be sufficient with low false-positive rates, data association can likely be improved by also considering inclination angle (easily incorporated for stabilized images),

range information (if reasonable estimates are obtainable by inclination calculations, target image size, or other sensors), time-history image similarities, and/or class information (when more contact classes are added). Ultimately, advancements such as these will aid in producing an omnidirectional maritime perception system capable of identifying, classifying, and tracking a variety of targets (on sea, on land, and in the air), enabling long-term, reliable ASV patrol operations.

ACKNOWLEDGMENTS

The research described in this paper was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration. Funding for this work was provided by the Office of Naval Research, Code 33 (Contract #N00014-09-IP-2-0008), and Spatial Integrated Systems, Inc. (NASA Space Act Agreement Contract #NMO716027). Finally, we wish to thank David Trotz, Robert Steele, Harry Balian, Lucas Scharenbroich, Mike Garrett, Lee Magnone, Tien-Hsin Chao, the NAVSEA boat drivers, SIS, and Seaward Systems for their valuable contributions to this project.

REFERENCES

- Benjamin, M., Curcio, J., Leonard, J., & Newman, P. (2006, May). Navigation of unmanned marine vehicles in accordance with the rules of the road. In *Proceedings IEEE International Conference on Robotics and Automation (ICRA)*, Orlando, FL.
- Benjamin, M. R. (2002a). Interval programming: A multi-objective optimization model for autonomous vehicle control. Ph.D. Thesis, Department of Computer Science, Brown University, Providence, RI.
- Benjamin, M. R. (2002b, October). Multi-objective autonomous vehicle navigation in the presence of cooperative and adversarial moving contacts. In *Proceedings of OCEANS*, Biloxi, MS.
- Chien, S., Knight, R., Stechert, A., Sherwood, R., & Rabideau, G. (2000, April). Using iterative repair to improve the responsiveness of planning and scheduling. In *Proceedings of the Fifth International Conference on Artificial Intelligence Planning and Scheduling (AIPS)*, Breckenridge, CO.
- Chien, S., Sherwood, R., Tran, D., Castano, R., Cichy, B., Davies, A., Rabideau, G., Tang, N., Burl, M., Mandl, D., Frye, S., Hengemihle, J., Agostino, J., Bote, R., Trout, B., Shulman, S., Ungar, S., Gaasbeck, J. V., Boyer, D., Griffin, M., Burke, H., Greeley, R., Doggett, T., Williams, K., Baker, V., & Dohm, J. (2003, May). Autonomous science on the EO-1 mission. In *Proceedings of the International Symposium on Artificial Intelligence, Robotics, and Automation in Space (i-SAIRAS)*, Nara, Japan.
- Hansen, E., Huntsberger, T., & Elkins, L. (2006, April). Autonomous maritime navigation: Developing autonomy skill sets for USVs. In *Proceedings of SPIE, Unmanned Systems Technology VIII*, Orlando, FL (pp. 221–232).
- Huntsberger, T., Cheng, Y., Baumgartner, E. T., Robinson, M., & Schenker, P. S. (2003, July). Sensory fusion for planetary surface robotic navigation, rendezvous, and manipulation operations. In *Proceedings International Conference on Advanced Robotics (ICAR)*, Coimbra, Portugal (pp. 1417–1424).
- Huntsberger, T., & Rose, J. (1998). BISMARC. *Neural Networks*, 11(7/8), 1497–1510.
- Huntsberger, T., Trebi-Ollennu, A., Aghazarian, H., Schenker, P. S., Pirjanian, P., & Nayar, H. D. (2004). Distributed control of multi-robot systems engaged in tightly coupled tasks. *Autonomous Robots*, 17, 79–92.
- Huntsberger, T., Trebi-Ollennu, P. P. A., Nayar, H. D., Aghazarian, H., Ganino, A., Garrett, M., Joshi, S. S., & Schenker, P. S. (2003). CAMPOUT: A control architecture for tightly coupled coordination of multi-robot systems for planetary surface exploration. *IEEE Transactions on Systems, Man and Cybernetics, Part A: Systems and Humans, Special Issue on Collective Intelligence*, 33(5), 550–559.
- Jimenez, A., Ceres, R., & Seco, F. (2004, June). A laser range-finder scanner system for precise manoeuvre and obstacle avoidance in maritime and inland navigation. In *Proceedings 46th International Symposium Electronics in Marine (ELMAR)*, Zadar, Croatia.
- Larson, J., Ebken, J., & Bruch, M. (2006, April). Autonomous navigation and obstacle avoidance for unmanned surface vehicles. In *SPIE Proceedings 6230: Unmanned Systems Technology VIII, Defense Security Symposium*, Orlando, FL.
- Nerveghna, M., & Ricard, M. (2006, August). Risk-aware mixed-initiative dynamic replanning (RMDR) program update. In *Proceedings AUVSI Unmanned Systems North America*, Orlando, FL.
- Perera, A., & Hoogs, A. (2004, August). Bayesian object-level change detection in grayscale imagery. In *Proceedings 17th International Conference on Pattern Recognition (ICPR)*, Cambridge, UK.
- Pirjanian, P. (1999). Behavior coordination mechanisms: State-of-the-art (Tech. Rep. IRIS-99-375). Institute for Robotics and Intelligent Systems, School of Engineering, University of Southern California.
- Pirjanian, P. (2000). Multiple objective behavior-based control. *Journal of Robotics and Autonomous Systems*, 31(1–2), 53–60.
- Saffiotti, A., Konolige, K., & Ruspini, E. H. (1995). A multivalued logic approach to integrating planning and control. *Artificial Intelligence*, 76, 481–526.
- Snyder, F. D., Morris, D. D., Haley, P. H., Collins, R., & Okerholm, A. M. (2004, October). Autonomous river navigation. In *Proceedings of SPIE, Mobile Robots XVII*, Philadelphia, PA (pp. 221–232).
- Yen, J., & Pfluger, N. (1995). A fuzzy logic based extension to Payton and Rosenblatt's command fusion method for mobile robot navigation. *IEEE Transactions on Systems, Man and Cybernetics*, 25(6), 971–978.