A Pilot-Aid for ROV Based Tracking of Gelatinous Animals in the Midwater

Jason Rife and Stephen M. Rock
Stanford University, CA 94305
and Monterey Bay Aquarium Research Institute, Moss Landing, CA, 95039
Email: jrife@arl.stanford.edu, rock@arl.stanford.edu

***Abstract*-A visual tracking system installed on MBARI ROV *Ventana* has demonstrated fully autonomous tracking of a gelatinous animal in the waters of Monterey Bay, California. The device, intended as an aid for human ROV pilots, will increase the feasible duration of observational experiments in the field of marine biology. This work describes the automation system and the challenges encountered during its development. These challenges have been addressed during ocean dives, throughout which the system has demonstrated robustness to the ROV operational environment and to the difficulties of imaging in the deep ocean.**

I. INTRODUCTION

Marine biologists benefit from remotely operated vehicles (ROVs) that gather detailed information about gelatinous animals in the ocean depths. Many deep ocean jellies are too fragile for transport, and must be studied *in situ*. The term "jellies" collectively refers to ocean dwelling gelatinous animals ranging in size and species. Sample gelatinous targets include 3-cm wide hydromedusae, 20-m long siphonophores, and 10-cm wide larvaceans. These animals are difficult to study using remote sensors mounted on surface vessels. Sonar waves pass through individual jellies with little reflection. Net trawls damage or obliterate fragile gelatinous specimens. By contrast, video captured by an ROV camera at short range provides useful insight into jelly behavior, diet, and life cycle.

Pilot fatigue places a limit on the duration of visual jellyfish studies. The jelly tracking task demands repetitive but precise thruster corrections. These continuous, fine-tuned maneuvers become stressful to a human pilot over time. For this reason, few attempts have been made to research a single specimen over an extended duration using a human piloted ROV.

Although current computer technology cannot match the ability of the human to adapt to new and unexpected situations, computers surpass human pilots in their ability to handle repetitive tasks with high precision. This discrepancy between computer and human capabilities gives rise to a spectrum of autonomy, with varying levels of human-robot interaction between pure human control and pure computer control. An increase in tracking duration implies a shift toward the automated end of the spectrum. A desire to perform occasional complex maneuvers or to react to unexpected scientific observation requires a shift back to the human end of the spectrum. Thus, the capabilities of humans and machines mesh to produce a flexible tracking system, with the potential for a tracking duration of multiple hours.

A jelly tracking pilot-aid has been implemented aboard the Monterey Bay Aquarium Research Institute (MBARI) ROV *Ventana*. From the earliest phases of the project, the tracking system has been tested in the open ocean in an attempt to understand the limitations and hurdles associated with field operations. Professional ROV pilots at MBARI have participated in the implementation, enhancing system applicability and ease-of-use. An early version of the jelly tracking system has demonstrated fully automated jelly tracking in the open ocean.

Fig. 1. MBARI ROV *Ventana*.

Other researchers have touched on the topic of automated animal tracking in natural underwater environments, but none has discussed real-time tracking of a gelatinous animal *in situ*. Minami *et al.* tackled robotic tracking of a fish, confined to a small tank [1]. Other workers have automated visual extraction of marine animals from a video sequence, without closing servo loops. Kocak *et al.* discuss vision techniques for off-line analysis of bioluminescent zooplankton data [2]. Fan and Balasuriya tested a 20 Hz fish tracking technique off-line, using video collected in the open ocean [3]. Other investigators have focused on pattern recognition methods useful for detecting underwater targets [4,5].

This work discusses experiments in ROV based, closed-loop tracking of jellies in the open ocean. Section 2 describes the experimental platform used to demonstrate fully automated jelly tracking. Section 3 describes some of the challenges in implementing an ROV pilot-aid. Section 4 outlines results of ocean experiments.

II. EXPERIMENTAL SYSTEM

A. ROV Platform

The jelly tracking system has been implemented on MBARI ROV *Ventana*. Video signals travel via a 2000 m fiber between the submerged robot and the surface support vessel, R/V *Point Lobos*. On board the support vessel, a 700 MHz Pentium III vision computer receives the video signal. The vision-processing algorithm segments candidate jelly targets from the video sequence and identifies the most likely candidate as the target. A PD controller holds the target on the optical axis of the ROV camera. Control commands, routed through the pilot joystick, are subject to a manual override to ensure safe operation of the submersible. The umbilical transmits these control commands, along with other telemetry and ROV power, between the support ship and the submersible.

B. Actuation and Control

Six thrusters actuate the vehicle in four degrees of freedom. Each of three thruster pairs permits control along a primary axis of the body reference frame, labeled with "b" subscripts in Fig. 2. A differential signal to the fore-aft thrusters accelerates the vehicle in the yaw coordinate. A passive buoyancy moment stabilizes the vehicle around the pitch and roll axes. ROV pitch and roll variations did not exceed 4° throughout experiments.

Cylindrical coordinates best suit the tracking problem in the absence of pitching and rolling motions. Fig. 3 depicts the tracking coordinate system from a top down viewpoint. Only three coordinates are necessary to describe jelly position in the body frame: range to jelly, r_b, yaw bearing, ψ, and relative depth, z_b. The fourth controllable degree of freedom, labeled n_b in Fig. 3, is redundant for the jelly tracking task. A more sophisticated system might take advantage of this null space for other objectives such as tether management.

Control design assumes a simple plant model with linear drag terms. Although bluff body drag models predict quadratic variation of drag with velocity at high Reynolds numbers (10^3-10^5), system identification for

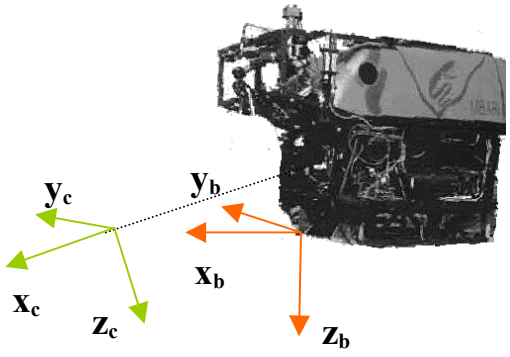

Fig. 2. ROV reference frames. The "b" subscript indicates the body-fixed frame, while the "c" subscript indicates the camera frame.

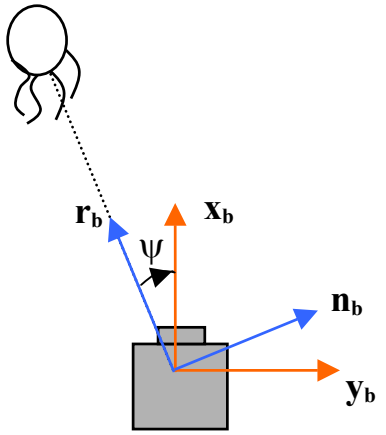

Fig. 3. Relationship between cylindrical tracking coordinates and Cartesian body frame coordinates. As the ROV and target are viewed from above, the z-axes, not shown, both point into the page.

Ventana suggests that a linear model adequately describes the system under jelly tracking conditions. The system drag model makes an additional approximation in expressing drag as a function of ROV motion relative to the jelly target, and not relative to the surrounding fluid. This simplified model works sufficiently well for the jelly tracking application and does not require the use of additional fluid motion sensors.

Given the assumptions of the linear drag model, neutral buoyancy and minimal tether forces, the dynamic equations for this six-state system are:

$$\begin{aligned} \mathbf{x} &= \begin{bmatrix} \dot{r}_b & \dot{z}_b & \dot{\psi} & r_b & z_b & \psi \end{bmatrix}^T \\ \mathbf{u} &= \begin{bmatrix} u_r & u_z & u_\psi & 0 & 0 & 0 \end{bmatrix}^T \\ \dot{\mathbf{x}} &= \begin{bmatrix} -\boldsymbol{\beta} & \mathbf{0} \\ \mathbf{I} & \mathbf{0} \end{bmatrix} \mathbf{x} + \mathbf{u} \end{aligned} \qquad (1)$$

The matrix $\boldsymbol{\beta}$ is diagonal, with entries equal to the damping coefficients for each axis normalized by the appropriate mass or inertia. Added mass is included in the $\boldsymbol{\beta}$ matrix. The control vector, $\mathbf{u}$, is written with units equal those of the time derivative of the state vector. The control axes map directly to the ROV thrusters, with the

exception of the radial control component. Radial thrust can be decomposed onto the body-frame Cartesian axes with the following rotational transformation:

$$\begin{bmatrix} u_x \\ u_y \end{bmatrix} = \begin{bmatrix} \cos(\psi) & -\sin(\psi) \\ \sin(\psi) & \cos(\psi) \end{bmatrix} \begin{bmatrix} u_r \\ u_n \end{bmatrix} \quad (2)$$

Since all axes are decoupled, classical control techniques may be used. Although several linear and nonlinear control techniques have been tested on *Ventana*, PD control performs well for the jelly tracking application. Equation (3) describes the control vector **u**, where $\mathbf{K}_d$ and $\mathbf{K}_p$ are 3x3 diagonal control gain matrices:

$$\mathbf{u} = -\begin{bmatrix} \mathbf{K}_d & \mathbf{K}_p \\ \mathbf{0} & \mathbf{0} \end{bmatrix} \mathbf{x} \quad (3)$$

C. Visual Sensing

Processing routines identify the jelly target in real-time within images captured by the main ROV camera, Sony model HDC-750a. The details of the image processing algorithm are reported in [6]. The location of the target centroid in a 2-D image provides a bearing measurement in the camera reference frame. Measurements of camera pan, tilt, and shoulder angles are required to transform from camera coordinates, denoted by a "c" subscript in Fig. 2, to body coordinates, denoted by the "b" subscript. Hall effect sensors located in the camera mount record angle measurements.

Range accompanies bearing to complete the description of target location in the body frame. Stereo triangulation, depth from focus, and feature based scaling are useful techniques for extracting range from an image sequence. Of these techniques, stereo triangulation and feature based scaling were tested.

Stereo triangulation uses a camera pair, with known separation and relative rotation, to calculate target coordinates. Fig. 4 depicts the general stereo geometry. With four bearing measurements used to calculate a three-element position vector, the problem is overconstrained. Kanatani suggests using a least squares approach, which results in the following equation [7]:

$$\rho_1 = \frac{\mathbf{t}^T \xi_1 - \xi_1^T \mathbf{R} \xi_2 \mathbf{t}^T \mathbf{R} \xi_2}{1 - (\xi_1^T \mathbf{R} \xi_2)^2} \quad (4)$$

Here ρ_1 is the distance to the jelly in the frame of the main camera, C_1. The unit vectors, $\boldsymbol{\xi}_i$, express the bearing to the target in the appropriate camera reference frame, C_i. **R** is the rotation matrix transforming the basis C_2 into C_1. Vector **t** is the translation between the two frames.

Feature based scaling produces a relative range measurement with arbitrary length scale. An external reference measurement is required to fix the units of the arbitrary length scale. Typically, *a priori* knowledge of target size or triangulation are used to relate a relative initial length to absolute metric units.

Relative length scales are based on pixel measurements in the image plane. Sample scales include square root of target area or target width. Pixel width is particularly appropriate for measurements of radially symmetric, approximately ellipsoidal targets, like most jellies. Rotational invariance holds for the minimum pixel width metric in the case of prolate targets, and for maximum pixel width in the case of oblate targets.

The experimental system used both triangulation and feature based scaling to acquire range measurements. In practice, the triangulation system, which employed a secondary camera mounted at a near right angle to the main camera, suffered from frequent dropouts when the target was visible in only one of the two camera images. For this reason, the code defaults to feature based scaling,

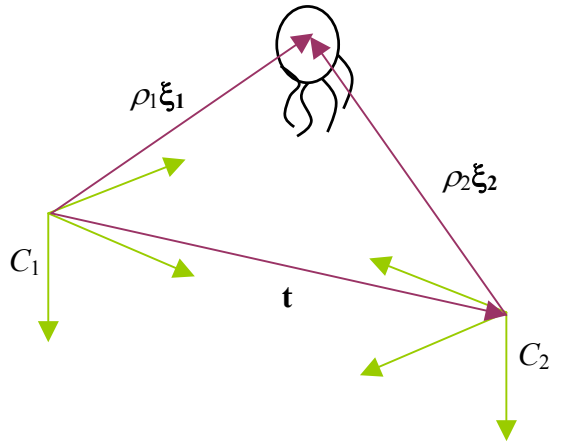

Fig. 4. Stereo triangulation geometry. C_1 and C_2 are camera reference frames, separated by the translation vector **t** and the rotation matrix **R**.

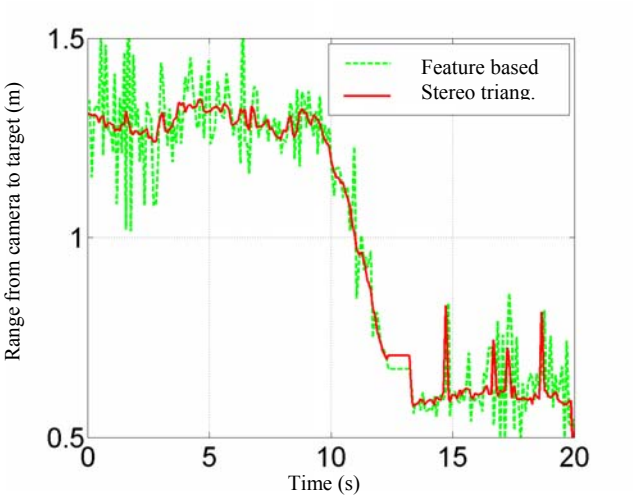

Fig. 5. Range to jelly as a function of time. The noiser signal was computed by feature based scaling. The cleaner signal resulted from stereo triangulation.

with initial range set at a pre-specified distance. Whenever two video streams become available, the system switches to triangulation ranging and recalibrates the initial range estimate for subsequent feature based scaling.

Fig. 5 shows raw data for simultaneous stereo triangulation and feature based scaling. The data was taken in the field, estimating range between the ROV camera and a target *Ctenophore*. Noise in the feature based range estimate results from difficulties in precisely determining target contours through segmentation. The feature based scaling signal can be smoothed, but at the cost of introducing a slight lag into the range control loop.

III. IMPLEMENTATION CHALLENGES

A. ROV Configuration

The open-frame construction of ROVs encourages frequent payload reconfiguration according to the requirements of each new task. Research ROVs see their payload change on a daily basis, as individual sensors or toolsleds are added to and removed from the submersible. At MBARI, the same ROV may be used to survey a clam field on the ocean bottom, to take core samples from a steep wall, and to conduct chemistry experiments in the midwater. Each system repair or reconfiguration introduces some likelihood of disrupting an unrelated component. The shear complexity of an instrumented ROV means that the overall probability of a system disruption is not trivial.

With a hybrid system combining human and computer supervision, the human operator can bear the burden of fault detection. A fully computer-controlled robot, by comparison, would require an automated fault detection routine. In order to evaluate system problems, the automated fault detection system would use redundant sensors or compare a system's dynamic behavior with a stored model. The human operator can more quickly and easily identify major system issues with a predive check and with continuous observation during ocean operations.

B. Lighting Environment

To locate a jelly in an image, a critical distinction must be made between pixels belonging to the target and those belonging to the background. At first glance, this segmentation process appears to be a trivial distinction between bright target pixels and dark ocean pixels. In fact, the presence of undesired features in the background, coupled with the transparent nature of most jellies, creates a segmentation challenge.

The dominant features of the background image are gradual intensity gradients created by scattered light. Scattering results as light travels from the illumination source and strikes microscopic particles suspended in the water. These redirected rays aggregate to form a diffuse light field that returns to the ROV camera. The fraction of light redirected by any differential volume of seawater depends on the relative angle between the incoming light ray and the receiving optics. The angular dependence of scattering, along with exponential absorption effects along the light path, is responsible for gradients in the light field arriving at the imaging array. Gradients characteristic of those observed in the field are depicted in Fig. 6.

Two phenomena alter these smooth, gradual gradients created by volumetric scattering. One effect results from reflection of a light source in the glass dome of the camera housing. These reflections appear as flares in the image. Flares are static, but bright, and may be mistaken as a jelly target. A second effect results from movement of large particles, known as marine snow, through the water column. The continual presence of marine snow adds small temporal and spatial variations to the otherwise smooth background image. The background pattern, modified slightly by snow motion, remains essentially unchanged for long periods of time, until the pilot or scientist requests a camera zoom, pan, tilt, or gain change.

Fig. 6. A *Ctenophore*, commonly referred to as a comb jelly, viewed in uneven background illumination. Contrast has been increased by a factor of four to enhance background variations caused by scattering.

In this work, a morphological gradient threshold has proved effective in segmenting jellies from a video sequence [6]. This filter is simple to apply and filters out most background lighting effects. Although the gradient filter does not eliminate marine snow, most particulates can be distinguished from the target by their smaller pixel area. The largest advantage of the gradient filter is its robustness, independent of camera zoom, camera gain, marine snow concentration, or lighting configuration. Unlike an intensity threshold, the gradient threshold need not be dynamically configured for different imaging environments.

Fig. 7 demonstrates the comparative effectiveness of threshold segmentation using intensity and morphological gradient thresholds. The figure depicts misclassification as a function of threshold choice, averaged over 45 samples. For a wide range of imaging conditions, the percentage of misclassified pixels remains low with gradient threshold fixed near 3-gray levels per pixel. Intensity based segmentation is not effective using a fixed threshold, as the best threshold changes substantially from image to image. Data for the plots draws 3 samples each from fifteen video sequences depicting different jelly species. The 45 samples were selected to emphasize imaging under

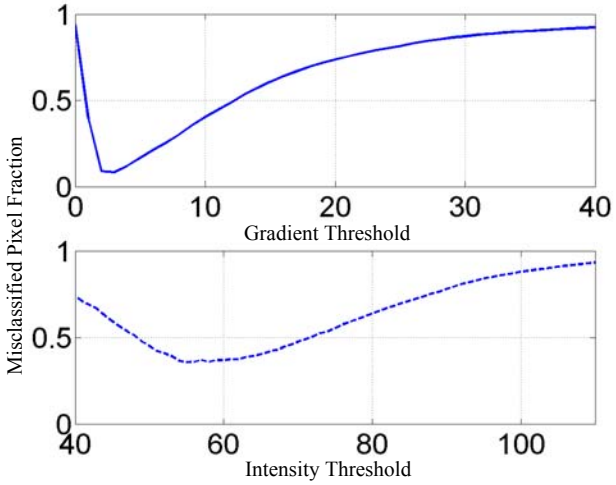

Fig. 7. Fraction of target pixels misclassified, as a function of threshold level. The upper graph (a) depicts misclassification for gradient based thresholding. The lower graph (b) depicts misclassification error for intensity based thresholding.

different lighting configurations with different camera settings. The misclassification metric was calculated at each threshold level as the maximum of two statistics, the percentage of background pixels misclassified as target pixels and the percentage of target pixels misclassified as background pixels.

C. Target Recognition

A visual tracking system should handle video clips containing more than one jelly as well as video clips containing no jellies. Fig. 8 depicts two of many instances involving multiple animals appearing in the camera frame. In the multiple-candidate case, the challenge resides in distinguishing the target among many segmented regions. Since segmentation techniques commonly identify flares and large debris particles as potential targets, the multiple-candidate case is actually the standard operational mode. Distinguishing between two gelatinous animals, however, requires finer resolution recognition than distinguishing between a jelly and a flare.

The recognition task is somewhat more challenging in the case of an out-of-frame event. At any time, a finite probability exists that the target does not appear in a video frame. In the ocean, sudden jelly acceleration, unplanned motion of the ROV caused by tether jerk, or local fluid motion may sweep the jelly outside the camera field of view. For this reason, the recognition routine must consider the possibility of a false positive match.

The experimental jelly tracking system uses a pattern vector to compare candidate regions to a target profile. The routine accepts as the target that region best matching the profile, as long as the match falls within false positive bounds. Both the pattern profile and the false positive boundaries are updated dynamically through time.

Fig. 8. Cases of multiple animals viewed in an image sequence. The left image (a) depicts a squid and a scyphomedusa. The right image (b) depicts two hydromedusae.

Recognition requires patterns that distinguish jellies from other animals and from each other. Most statistics describing the target vary with time or position in the image. Lighting gradients and spectral variations in absorption, for instance, imply that image intensity and color are not invariant to jelly position. Recent position and velocity estimates help distinguish among candidates, but not uniquely or continuously. Target trajectory fails as an identification tool when two jellies slide closely by each other or when the target jelly leaves and re-enters the field of view. Even target size and shape, projected into the two-dimensional image plane, vary significantly with jelly rotation and deformation. Tentacled jellies, in particular, employ different body configurations for resting, swimming and hunting. The three frames of Fig. 9, snapped at one-second intervals, show *Colobonema*

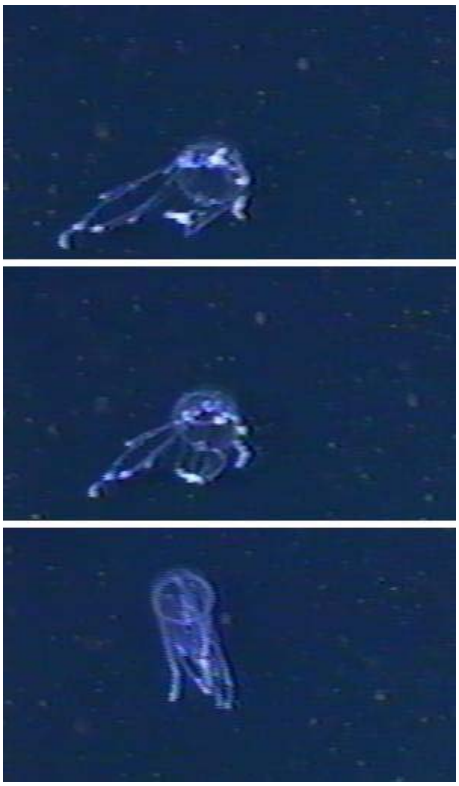

Fig. 9. Transition from rest (top) to motion (bottom). The tentacles of this *Colobonema* are significantly rearranged when the jelly accelerates.

accelerating from a stationary configuration into a swimming configuration.

The experimental system combines all these statistical descriptors together in an attempt to recognize the jelly target [6]. The recognition routines performed well in field trials, where the appearance of multiple animals in the camera frame was a frequent occurrence. The system did, however, occasionally misclassify the true target as a false positive. The frequency of incorrect declaration of a false positive depends on the species, its range from the camera, and the sensitivity of the false positive boundary. The human pilot provides a safety net, here, deactivating the automated tracking system during loss of lock, and reactivating the system a short time later.

D. Target Initialization

The jelly tracking task, well suited to computer control, requires a relatively complex initialization, better suited to human control. Three steps define the initialization procedure. First, the ROV enters a search mode, during which it follows a transect trajectory while looking for potential targets inside the field of rapidly passing marine snow. Observation of a potential target triggers a second, deceleration mode. In this mode, the vehicle must come to rest relative to the water column without losing sight of the perceived target. The ROV subsequently enters a recognition mode, during which the vehicle hovers near the potential target in an attempt to confirm the target's identity. If the species of the potential target matches the species required for a scientific experiment, control transitions to the automated jelly tracking algorithm. Otherwise, the vehicle must return to the search mode and repeat the process.

Humans are effective at performing all three initialization steps. Pilots and scientists have a particular advantage over computers in identifying the species of an animal, especially one that is deformable and is viewed in an unstructured environment. This class of difficult recognition task continues to excite the computer vision community as a topic for research. Given the current state of the art in visual recognition, the hybrid system works more effectively than a pure computer-supervised solution in that it places the initial identification task in human hands.

A simple interface is a practical requirement for the jelly tracking system. Until computer tracking begins, the pilot is tightly focused on maintaining the target jelly in the camera field of view. Under these stressful conditions, a single button interface proves very effective. For this reason, the experimental system permits activation with the press of a signal button. The computer controller regulates ROV position around a set of coordinates established at the time of the button press. This simple interface has precedent in existing pilot navigation aids, such as auto-depth and auto-heading systems.

As part of the single button interface, the automated tracking system must select an initial target from the video stream and train a profile for that target. In most cases, this task is made easier by leveraging the human component in the hybrid control system. For experiments, the human pilot was asked to align the target jelly with the camera optical axis and ensure only a single jelly appeared in the video sequence. Using these assumptions, the automated jelly tracker rarely misidentified the initial target.

Initialization of the automated system should persist long enough to build a recognition profile for the target specimen. By initializing the profile at run-time, rather than in advance, the tracking system retains simplicity and a high degree of flexibility to track any species encountered. This capability is significant, as scientists continue to discover new species, adding to the list of thousands of gelatinous animals already identified. The visual system may require the human pilot to retain control during this learning phase, in which case the learning phase should remain short. The training phase for the experimental system lasted for less than 2 seconds.

E. Scientist Interaction

A scientist will frequently choose to interact with the jelly tracking system. Common interactions include changes to camera tilt, pan, zoom, or gain settings. When the autopilot system shares a camera with the scientist, these changes may profoundly affect automated tracking. Camera changes can directly modify the character of the background image, as described in Section B. Zoom can also change the patterns that identify the target profile, as described in Section C. The transfer function between image plane coordinates and bearing angle also change with zoom, effectively altering hard-coded control gains.

In some cases, the scientist may desire that the vehicle perform a complex maneuver, such as a rotation around the target jelly to view it from a new angle. Since human pilots respond well to complex and unplanned demands, these situations are best handled by switching from computer to human control. In this sense, the hybrid human-computer control scheme offers significant mission planning flexibility to a research scientist.

IV. RESULTS

During July 2001, *Venatana* accomplished the jelly tracking task under full computer control. This initial attempt at jelly tracking lasted for 2.5 minute. Task difficulty was enhanced by the small size of the target, which demanded tracking at a short range, approximately 0.5 m. High sensitivity at close range increased the likelihood of losing the jelly from the image frame. Fig. 12 depicts position and control signals along all three tracking coordinates during the 2.5-minute period of fully automated jelly tracking. Note the ROV was displaced from its vertical reference by approximately 0.1 m due to a constant disturbance load (Fig. 12a).

Longer tracking runs have been demonstrated, but under partial computer control. An earlier test at a range of 3 m lasted for 10 minutes. This test benefited from lower controller sensitivity at a longer range from the

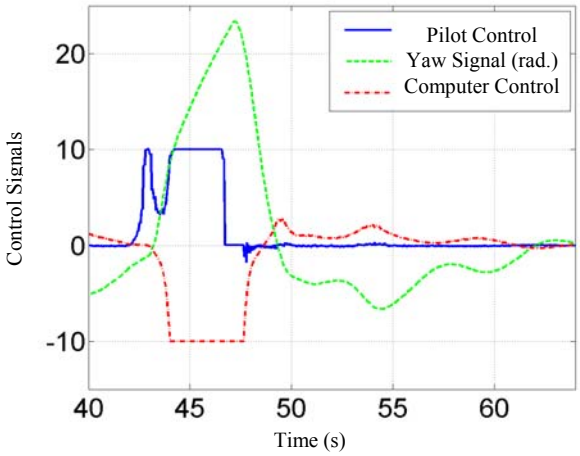

Fig. 10. Control signal for the jelly tracker during disturbance rejection. The image sequence of Fig. 11 runs from 41 seconds to 50 seconds.

target. During these trials, the pilot controlled range while the computer controlled yaw and depth. After tracking the target jelly for 10 minutes, the algorithm, which did not yet include recognition capability, failed when a squid entered the scene (Fig. 8a). The pilot corrected the system after the squid had disappeared, and tracking continued for another ten minutes.

Fig. 11 shows the control loop's ability to reject disturbances. The video sequence shows ten successive

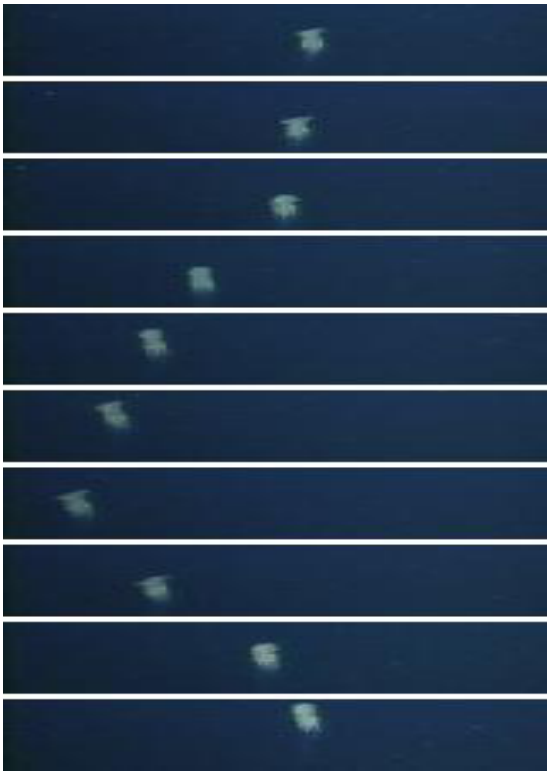

Fig. 11. Disturbance rejection by the jelly tracking control loop. Images are sampled with at a one second interval.

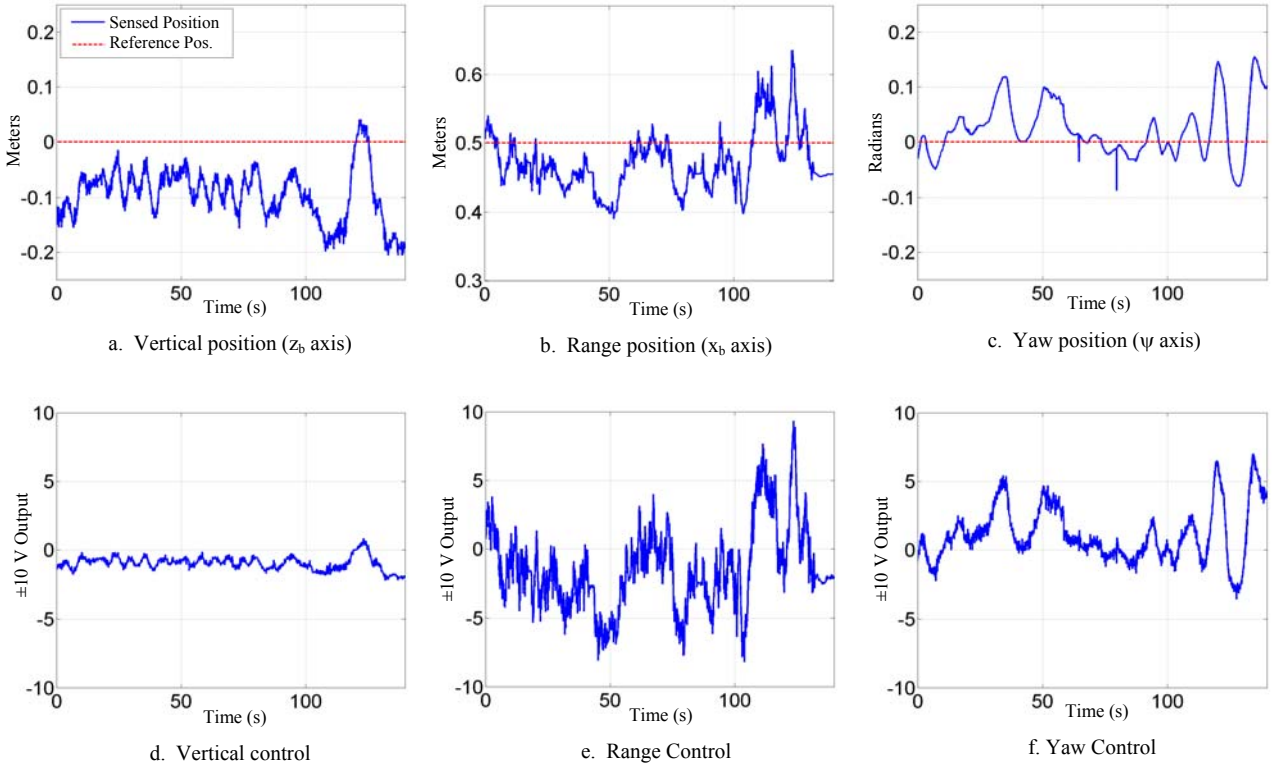

a. Vertical position (z_b axis)

b. Range position (x_b axis)

c. Yaw position (ψ axis)

d. Vertical control

e. Range Control

f. Yaw Control

Fig. 12. Sensor measurements (a-c) and control signals (d-f) during 2.5 minutes of fully autonomous jelly tracking.

images of an egg-yolk jelly, *Phacellophora camtschatica,* sampled at one-second intervals. During the first five images of the series, the human pilot gradually rotates the ROV *Ventana* around the yaw axis. During the subsequent five images, the automatic controller snaps the jelly back to the neutral yaw position. Fig. 10 depicts control forces applied by the human pilot and by the automatic controller around the yaw axis. The image sequence of Fig. 11 maps to the time axis of Fig. 10, between 41 and 50 seconds.

V. FUTURE WORK

Short-term activities will focus on improvements to system identification and control that should lengthen the duration of fully automated jelly tracking. Vision routines will also be improved, to permit better recognition and false positive detection. Future efforts will examine automation of some of the system initialization tasks, described in Section III.D. Power requirements for the sensor will also be addressed, in an effort to produce a system practical for AUV applications.

VI. CONCLUSION

An automated jelly tracking system was implemented and tested on ROV *Ventana* in the open ocean. By placing repetitive, precision maneuvering under computer control, the jelly tracker frees human pilot attention to perform higher level decisions and to react to unplanned situations. The system successfully tracked a target jelly for 2.5 minutes under full computer control and for as long as ten minutes with primary computer control, tweaked by pilot input in one degree of freedom. The visual jelly tracking sensor was designed in the context of scientific ROV operations, with provisions for the particular challenges presented by this operational environment including frequent ROV reconfiguration, scattering of light in the ocean, and the appearance of other animals near the tracked specimen. The jelly tracker was designed to initialize with a single button press, so that the pilot interface remains simple.

ACKNOWLEDGEMENTS

We thank MBARI and Packard Foundation grants 98-3816 and 98-6228 for supporting this work.

REFERENCES

[1] M. Minami, J. Agbanhan, T. Asakura. "Manipulator Visual Servoing and Tracking of Fish Using a Genetic Algorithm", *Industrial Robot*, v. 26, 4, pp. 278-289, 1999.

[2] D. Kocak, N. da Vitoria Lobo, E. Widder. "Computer Vision Techniques for Quantifying, Tracking, and Identifying Bioluminescent Plankton", *IEEE Journal of Oceanic Engineering*, 24, v. 1, pp. 81-95, 1999.

[3] Y. Fan and A. Balasuriya. "Autonomous Target Tracking by AUVs using Dynamic Vision", *Proc. of the 2000 International Symposium on Underwater Technology.* pp. 187-192, 2000.

[4] X. Tang and W. K. Stewart. "Plankton Image Classification Using Novel Parallel-Training Learning Vector Quantization Network", *Proc. IEEE/MTS OCEANS '96.* v.*3,* pp. 1227-1236. 1996.

[5] X. Yuan, Z. Hu, J. Chen, R. Chen, P. Liu. "Online Learning and Object Recognition for AUV Optical Vision", *IEEE Int. Conf. on Systems, Man, and Cybernetics, 1999.* v.6, pp. 857-862. 1999.

[6] J. Rife, S. Rock. "Visual Tracking of Jellyfish *in Situ*," *Proc. Intl. Conf. Image Processing*, IEEE, 2001, in press.

[7] K. Kanatani. *Geometric Computation for Machine Vision.* Clarendon, 1993.