Project Manager: Duane Edgington

900260: Visual Image Event Detection Technology Exploration

Project Manager: Duane Edgington
Lead Scientist: Bruce Robison
Lead Engineer: Duane Edgington
Duration: 2002

Date Submitted: August 20, 2001

1) Abstract

In this project we propose to investigate the means to visually process images for event detection and for the recognition of target biological species, with these systems to be deployed on autonomous underwater robotic vehicles (AUVs). This is a technology exploration project, to review both state of the art and state of commercial practice in detecting events (new features in the visual field) using video cameras, digital still cameras (CCD pixel arrays) and neuromorphic artificial retinas and photoreceptor arrays.

We seek to develop vision systems to add to AUVs to detect when they pass an organism, and to take an image of the organism. Our dream is to be able to identify the organism autonomously in real time.

We will evaluate application of feature-detection video frame grab technology developed by Steve Rock's group (600034 2001) being adapted to Tiburon. We will evaluate underwater and commercial digital still cameras, with illumination by strobe, as a potential lower power alternative to continuous light and video on AUVs. We will also evaluate state of the art feature and motion detection chips modeled after biological vision systems. These neuromorphic systems feature ultra-low power, large dynamic range and intrinsic real-time image processing and feature detection. Our application focus will include both midwater and deep ocean animals as well as animals that use bioluminescence in the deep sea.

2) Proposal

a) Statement of the problem:

One of several areas where unmanned remotely operated underwater vehicles (ROVs) have proven to be useful tools in oceanography is the capture of video and visual images of biological creatures. The great success of ROVs in collecting high-resolution video during dives encourages us to look for ways to increase the amount of information and knowledge we gather from underwater cameras by attempting to automate portions of the analysis of video, and to extend image collection to unmanned autonomous underwater vehicles (AUVs).

ROVs are natural platforms for collecting broadcast quality high-resolution video, since high-speed telemetry from vehicle to surface and electrical power for camera and lights are available via a cable tether. Some analysis or adaptive sampling is possible during the dive via surface video monitors and annotation systems, and outline or detailed post-dive analysis is routinely done in the video lab using sophisticated tools such as the MBARI Video Annotation and Reference System (VARS). There are two bottlenecks in collecting and analyzing this data that we hope to address in this project:

1. Video annotation is demanding and tedious work; highly trained scientists can work at most 4 hours a day annotating video; the capacity of the video lab and research staff to

post-process dive footage is a historical bottleneck in reaping the benefit of these expensive dives. One of our goals is to develop systems and method to augment the productivity of video annotation staff who are processing video from ROV dives.

2. The ROVs are expensive, high demand resources. The science staff would like to double or more the length of dives that are taking quantitative species and abundance samples (transects), as well as increase the frequency and sites of such transects, which is impracticable with the ROV resources and expense that MBARI has. It is very attractive to augment ROV dives with AUVs, which promise the capacity to go longer, more frequently, at lower cost than ROVs. However, ROVs are very power limited, since they are not tethered, and must carry their power sources and data storage with them during missions. One of our major goals is to investigate technology for enabling AUVs to fulfill the same kind of science missions that ROVs currently provide for species and abundance sampling.

The goal of this project is to explore high-risk high-payoff technologies that will enable detection of a visual event in the presence of "noise" in the visual field.   We will investigate both research work and commercial systems that do visual image processing and image capture.

This exploration project will be done with participation and input from

a) potential science users (Robison group and others),
b) video lab personnel and science researchers who analyze video (expertise on features of interest, "noise" in the visual field, and illumination requirements),
c) 600034 Steve Rock's group (expertise on automated video processing), outside organizations deploying cameras on AUVs, and the 600034 project team.
d) research organizations developing neuromorphic vision chips, and
e) 900250 "Quantitative Low-Light Imaging Module for MBARI's AUV" (Reisenbichler and Mellinger) to leverage their investigation of appropriate low power and low light camera systems for AUVs (expertise on underwater camera and lighting mechanical housing, illumination system engineering, electrical/electronic considerations).

Why it is important:

MBARI has significant investment in ROV video image capture technology and scientific investigation tied to video analysis. This project seeks technology that would augment the effectiveness of the staff and equipment to solve scientific questions. In addition, technologies to add autonomous vehicles and cameras to the scientific arsenal could prove extremely valuable in increasing our ability to sample more thoroughly, more often and in more places economically. The ability to autonomously detect and identify a creature could open up new lines of scientific enquiry, like the creature-following AUV proposed in Steve Rock's 600034 proposal. Finally, investigation of alternative image processing and detection techniques could suggest follow-on innovative work, such as optical-flow based navigation, feature following, and so on.

Present state of art/understanding:

Visual image processing is a huge field with a rich history of investigation. Automatic visual feature detection has been a research and engineering topic in a number of disciplines, including medical image analysis, surveillance, robotics (for example, place recognition) and industrial control. We won't try to summarize this vast literature here.

There are two classes of approaches to visual feature detection. The first is a parametric approach, where a captured image is processed to detect edges, lines, corners, junctions and so on, and then these parameters are used to match some model of the feature of interest. The images and parameters can be 2D, or 3D images can be constructed and analyzed. Video image processing can add multi-frame images to the analysis, comprehending blurring, rotation, noise, and shifts between frames. This enormous field of research and engineering has yielded surprisingly few robust general systems, but there are systems of focused application that are successful in feature detection and visual event detection. It is a major part of this project to sift through the work in this area to look for demonstrated systems and approaches that apply to our problem domain.

Another kind of approach to visual feature detection is neuromorphic engineering: silicon systems, following the metaphor of biological nervous systems, which process visual signals at a silicon photoreceptor array in parallel, rather than post-processing an image. This relatively recent approach has yielded interesting engineering systems, which like our own nervous systems are low power, highly parallel and scalable, reliable in the face of noise, and robust. The simple idea behind this approach is that biological systems have established the existence proof that visual events can be detected in a noisy environment, and features and creatures can be recognized because a trained human can do it. These systems seek to follow the guiding biological metaphor to build silicon subsystems using conventional VLSI techniques. Edgington started evaluating this technology at the 2001 Telluride Neuromorphic Engineering workshop, sponsored by the NSF and ONR, and generated interest in this approach to our problem domain from contacts at several of the significant research groups in neuromorphic engineering.

How it relates/contributes to MBARI's mission and strategic plan:

This proposal attempts to incorporate leading-edge research and engineering approaches to visual image processing to problems in oceanography, which to our knowledge is novel and new. The project if successful would offer new approaches to analyzing ROV images and the use of AUVs for species and abundance surveys that would have wide applicability and interest to oceanography. This high-risk, high-reward project would be unlikely to be undertaken at other oceanographic institutions, but would yield approaches of great interest to those researchers and engineers if successful.

b) Approach to the problem

Our program approach is to investigate specific algorithms or methods for detecting visual events, documenting those techniques, and then simulating or prototyping the most promising techniques. We propose no deployable implementations in 2002; if appropriate those would be proposed separately as follow-on projects to this proposal.

Project tasks are:

1. Investigate approaches for detecting a visual object in midwater in the presence of variable amount of marine snow.

2. Investigate algorithms for motion detection.

3. Investigate algorithms for “large object” detection, and identification based on characteristic features.

4. Develop simulation or prototype detection system based on most promising techniques and algorithms.

5. If appropriate, develop proof-of-concept implementation of visual image event detection system.

Each of these tasks is described below:

Task 1: Investigate approaches for detecting a visual object in midwater in the presence of variable amount of marine snow.

This is one of the main challenges that we perceive for the project – to pick out visual events in the visually noisy ocean environment of marine snow. One approach is to develop algorithms (which map to circuits or software) to detect optical flow of marine snow and report deviation from that flow of particles, which could signal a creature moving counter to the optical flow. This perceived “counter movement” often accompanies a swimming creature or an object much larger than the snow. Another approach is the more traditional feature analysis of a static visual image, classifying particles as snow based on some set of criteria, and reporting features not classified as snow to be potential events. It might well prove that a collection or hybrid of approaches is most effective.

Task 2: Investigate algorithms for motion detection.

There are a variety of motion detection algorithms, which build on task 1. “Bug detectors” are metaphors for motion detection in biological visual systems, and these systems are relatively well understood. Another neuromorphic motion detector is the Barlow-Levick motion selective silicon retina (Benson and Delbruck, 1997), which could be simulated and tested with our video collection. Other algorithms detect motion by noting blurring, change in focal distance, or change in position frame to frame that deviate from predicted position based on the velocity of the vehicle and optical parameters.

Task 3: Investigate algorithms for “large object” detection, and identification based on characteristic features.

“Large object” detection (larger than marine snow) builds on task 1 and 2. Foveation and saccadic scan are metaphors of how biological visual systems detect and track an object of interest in the visual field. More traditional algorithms are based on parametric feature identification, such as edges, texture, color, brightness, bioluminescence and characteristic subsystems of parameters (aspect ratio, limbs, eyes, etc.).

Task 4: Develop simulation or prototype detection system based on most promising techniques and hardware

We should be able to evaluate promising techniques by simulating them in software or building hardware prototypes in the lab incorporating the algorithms. The frame capture and analysis environment that 600034 has built up, and which Mike Risi is quite familiar with, could be a valuable test-bed for traditional visual image processing techniques. Matlab image processing toolbox 3 is another prototyping environment that could prove useful to test algorithms. A simulation environment for neuromorphic approaches is IQR421, available from the Institute of Neuroinformatics at the Swiss Federal Institute of Technology (ETH Zurich). Silicon algorithms can be prototyped economically by implementation in Tanner Research L-edit layout editor, using cell libraries freely available from the University of Pennsylvania Neuroengineering Research Lab and fabrication at Agilent/HP using the MOSIS research VLSI fabrication service. Research labs have implemented many silicon algorithms, and those circuits are usually available for research evaluation at cost. We have requested budget for some limited simulation tools and prototyping hardware.

Simulations or prototypes can be driven by stock video footage, selected in collaboration with the video lab personnel or research scientists to emulate typical dives and environments of interest.

Task 5: If appropriate develop proof-of-concept implementation of detection system.

Finally, if sufficiently promising approaches prove attractive in simulation or prototyping, we will attempt proof-of-concept implementations on appropriate hardware and software systems to demonstrate capability and explore robustness.

Relationship to other proposed efforts:

There are other projects being proposed for 2002 that are very significant in the context of our proposed project. Discussions have been held, and we all agree that it would be very helpful to all if we collaborate in our activities where appropriate.

900250 “Quantitative Low-Light Imaging Module for MBARI's AUV” (Reisenbichler and Mellinger) to leverage their investigation of appropriate low power and low light camera systems for AUVs (expertise on underwater camera and lighting mechanical housing, illumination system engineering, electrical/electronic considerations). This very pragmatic project, proposed in 2002 to evaluate technology, provides us with insight into how to mount a camera on an AUV, what the power budget is for camera and image storage, and low light CCD video and still camera options for AUVs. There is a great deal of synergy in each of our project’s need to understand lighting options and constraints on AUVs, and options for low-light continuous video capture.

600034 “Precision Control Technologies for ROVs and intervention AUVs” (Rock and McEwen) to leverage their on-going work in image capture and processing to detect features for control system “bottom-lock”, as well as jelly tracking (expertise on automated image processing). This project has developed working proven technologies and approaches that are very relevant to our proposal. There is a great deal of synergy in our joint interest in eventually developing capability to identify autonomously the presence of a particular creature.

900202 “Benthic image device” (Malby and Buck) who have a different problem domain (ocean bottom rather than midwater) which could prove to be a more tractable short-term objective than our long-range hope of identifying midwater animals.

c) Specific deliverables (2002)

Products from research (tools, publications, presentations, etc.)

The deliverables for this project will include:

1) White paper document summarizing and comparing available and emerging visual image processing techniques applicable to this application domain.
2) Proof-of-concept prototype(s) or simulations of one or more of the most promising techniques and hardware.
3) Follow-on project proposal and plan for implementation, test and deployment of visual image processing system.

How success will be measured

The ultimate measure of the success of our program is whether follow-on project proposals are generated and approved – i.e. can we successfully glean from the research literature and commercial systems approaches that are compelling and demonstrable.

A second measure is our demonstrated impact on related 2002 projects, enumerated above.

d) Project team

Who will be working on this project?

From MBARI: Duane Edgington (Lead Engineer), Bruce Robison (Lead Scientist), Mike Risi (Software Engineer), Nancy Jacobsen and other members of the MBARI video lab, and collaborators from other projects with whom we share common goals.

What each team member contributes to the solution of the problem:

Duane Edgington will be the project manager of the activities and will be lead engineer.

Mike Risi will be a software engineer participating in prototype development, and linkage to the 600034 project.

Bruce Robison (Lead Scientist) will provide direction and feedback on approaches and objectives.

d) Project Lab Space

We will do our prototype development work in the Software Engineering Lab.

3) Budget

Task 1: Investigate approaches for detecting a visual object in midwater in the presence of variable amount of marine snow.

Engineering:	Duane Edgington, Mike Risi	SE 30 days
Video Lab		5 days

Task 2: Investigate algorithms for motion detection.

Engineering:	Duane Edgington, Mike Risi	SE 30 days
Video Lab		5 days

Task 3: Investigate algorithms for “large object” detection, and identification based on characteristic features.

Engineering:	Duane Edgington, Mike Risi	SE 30 days
Video Lab		5 days

Task 4: Develop simulation or prototype detection system based on most promising techniques and hardware.

Engineering:	Duane Edgington, Mike Risi	SE 40 days
Video Lab		5 days
Equipment:	Lab computer	$5K
	IQR421	0

Lab Digital Camera	$1K
Matlab image toolbox	$1.35K
Tanner Research L-Edit	$5K
MOSIS Agilent/HP VLSI fab	$5.5K

Task 5: If appropriate develop proof-of-concept implementation of detection system

Engineering	Duane Edgington, Mike Risi	30 days

Additional budget:

Travel to Neuromorphic Engineering Workshop, 2002.	$3K
Travel to Image Processing Conference, 2002	$2K

4) Milestones

Milestones for 2002:

1) White paper document summarizing and comparing available and emerging visual image processing techniques applicable to this application domain. Q3 2002
2) Proof-of-concept prototype(s) or simulations of one or more of the most promising techniques and hardware. Q4 2002
3) Follow-on project proposal and plan for implementation, test and deployment of visual image processing system. Q3 2002

5) Curriculum Vitae

CVs on file.