NAVAL RESEARCH
REVIEWS
APRIL 1976

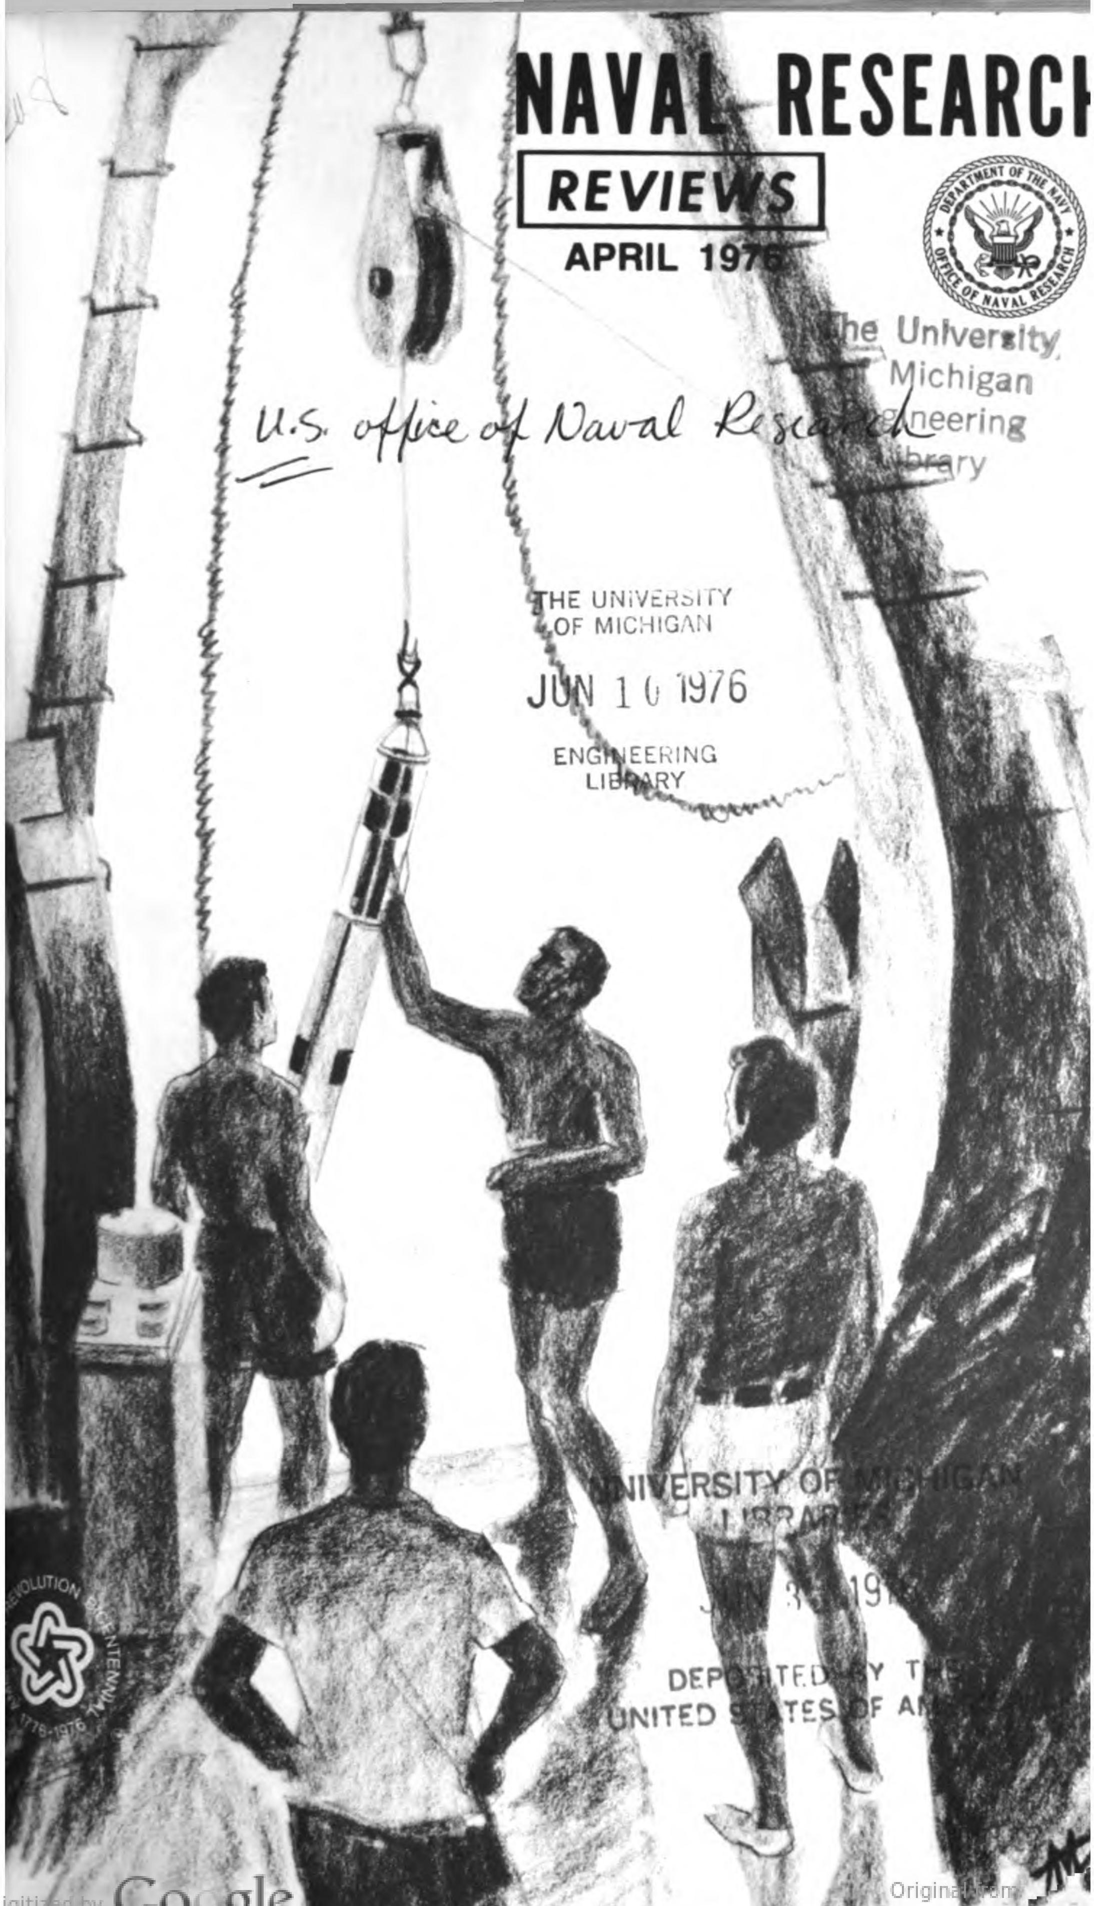

U.S. office of Naval Research

PROFILES IN SCIENCE

Professor Joseph B. Keller

Professor Joseph B. Keller is an applied mathematician. He is Director of the Division of Wave Propagation and Applied Mathematics at the Courant Institute of Mathematical Sciences, New York University. His research has concerned electromagnetic and acoustic wave propagation, fluid dynamics and other areas of applied mathematics. One of the results of his work is the Geometrical Theory of Diffraction, a widely used theory of the propagation and scattering of radar and sonar waves. Another is a method for the analysis of waves in random media, such as sound waves in a turbulent ocean and radio waves in a turbulent atmosphere. During 1953-4 he served as Head of the Mathematics Branch of ONR. He is on the editorial boards of various mathematics and physics journals, is a member of the National Academy of Sciences, and of the U.S. National Committee on Applied Mechanics, and is on the Board of Trustees of Society for Industrial and Applied Mathematics (SIAM).

Guidelines for Research

In order to establish within the Department of the Navy a consistent policy for research programs, Secretary of the Navy Middendorf sent the following memorandum to the Chief of Naval Operations and the Commandant of the Marine Corps.

In 1945, Secretary of the Navy James Forrestal addressed research in his report to the President: "I have said that it is even more important in peace than it is in war for scientists to contribute to basic military research. Wars are fought primarily with weapons which were developed before the fighting began...Research requires time. Application of its results absorbs more time.... If a weapon is to be employed during a war, the fundamental research underlying it must have been conducted over a protracted period before the war...."

Today, shortly after the close of a different kind of engagement in Southeast Asia, these fundamentals take on an even greater urgency than at the close of World War II. The political, economic, and social environment in which we live makes our job, more than ever before, one of assuring peace or minimum military response to those events which require it. Our true enemy is ignorance. A nation that must live in ignorance must live in fear. It is likely to respond inaccurately and over-aggressively to imagined threats and to overlook more subtle and pervasive developments. Accordingly, I request that you consider the following as statements of policy in developing programs and budgets in fundamental research as a part of our overall emphasis on strengthening and rebuilding the technology base:

a. Research is an important part of the overall Navy and Marine Corps Research and Development Program, and the quality and adequacy of size of the research effort are of utmost importance.

b. Since research is an investment recognized as long-term, the Navy budget for research must be protected from incursions prompted by current deficiencies in development and production programs.

c. The disciplines and programs supported by naval research should, in general, be those judged to be of greatest future potential for the needs of the Navy and Marine Corps. They should complement or strengthen the national research base in these areas. In the light of the Navy's particular or unique dependence on areas such as acoustic oceanography, our commitments to those disciplines should remain substantial.

d. One emphasis of the naval research program should be preservation and enhancement through the contract program of university, institutional and industrial capability to provide insight into the basic sciences and to interpret their potential for the Navy. In the long run, the Navy must help to assure stability in the critical base of academic research, so that the competence and independence of view of this community can contribute to a proper balance among the Navy's research programs, as well as to provide advice and counsel on matters critical to the Navy's future. It is'important for Navy/Marine Corps officiers to use this program as a means to expand their personal perspective through initiatives to meet and exchange ideas with these scientists.

e. The naval research program, while retaining its integrity and independence, should support our total technology base efforts and be responsive to research needs of developmental and operational programs. It should exercise leverage wherever possible through cooperative and synergistic efforts with the research and development programs of the Department of Defense, other government agencies, and allied foreign nations.

f. It is essential that the research program be characterized by adaptability and flexibility in order to follow up and assess the potential of new ideas and discoveries as rapidly as possible.

Further, in allocating research funds, I desire that our major emphasis of naval research continue to be oriented toward our constructive interaction with university research. Every effort must be made to minimize fluctuations in program level consistent with academic procedures and to assure an adequate influx of new investigators. Research must continue to command a proper priority in Navy programs.

Experimental Cold Oil Facility

At the Naval Coastal Systems Laboratory, Panama City Florida, there has been created an experimental facility that comes as close to an oil salvage situation in the Northern latitude as has been accomplished anywhere.

It's all part of NCSL's developmental program for evaluating equipment and techniques for pumping cold, highly viscous oils found aboard striken oceanic vessels.

After a series of tests, following installation of special equipments, the experimental facility is capable of pumping oil as low as 38 degrees Fahrenheit with a viscosity of 137,000 SSU (Saybolt Seconds Universal). At 39 degrees the oil is as thick as molasses or rubber cement. During the series of tests, oil was pumped at temperatures ranging up to 95 degrees Fahrenheit, or 135,000 SSU.

(Continued on Page 17)

Representation and Use of Knowledge in Vision

H. G. Barrow and J. M. Tenenbaum*
Stanford Research Institute

Compared to the rest of the computer field, the methods used in computerized Scene Analysis are very primitive. The Office of Naval Research recognized this defficiency and supported research to define the present state of Scene Analysis and develop goals to guide future research. The Navy will benefit from this work because it will lead to better automated systems in such areas as target detection and identification and unmanned vehicle control.

Introduction

Visual perception, whether by a human or a machine, is a process of explaining sensory data in terms of a priori models of the world. In its simplest form, the perceptual process might involve partitioning a picture into "coherent" regions, based on, say, models of texture or color homogeneity. Higher levels of perception could involve partitioning the picture into "meaningful" regions, based on models of particular objects, classes of objects, likely events in the world, likely configurations, and even on non-visual events. Vision might be viewed as a vast, multi-level optimization problem, involving a search for the best interpretation simultaneously over all levels of knowledge.

There is an important point to note here. The perceptual process, as applied to a given scene, does not have a single "correct" solution; the desired optimum must depend on the goals, interests, background and even sensory acuity of the perceiver.

This paper addresses a number of key AI (artificial intelligence) issues:

1) What knowledge is necessary to segment and interpret various classes of scene?
2) How should knowledge be represented and used?
3) Is there knowledge general enough to encompass a wide class of scene types?
4) How should knowledge be acquired?

*Dr. Barrow and Dr. Tennenbaum are Senior Research Engineers at Stanford Research Institute, Menlo Park, California. Their field of interest is artificial intelligence concentrating on computer perception.

Background

Research to date has resulted in a few scene analysis systems that perform tolerably well in segmenting and interpreting images in some limited domains, notably polyhedra, human faces, room scenes and simple landscapes, and a few uniform procedures for partitioning arbitrary pictures into coherent regions. The following table presents a few representative systems that indicate the current level of competence in the field.

The programs described in Table I depend strongly on characteristics of their particular domains. The Toy Worlds are usually characterized by sensorily deprived images (e.g. no shadows or texture), a small number of object prototypes, precise structural object descriptions,

TABLE I

I. TOY WORLDS.		
1: Polyhedra.		
	a Roberts	Interpretation of extracted line drawings as three dimensional objects, based on generic prototypes of wedge, rectangular prism and hexagonal prism [1].
	b Falk	Interpretation of extracted noisy line drawings based on nine fixed-size object prototypes [2].
	c Shirai	Extraction of line drawings for arbitrary configurations of bricks using heuristics that predict additional lines in a partially completed drawing [3].
	d Waltz	Interpretation of ideal line drawings as arbitrary trihedral solids based on a catalog of legal edge and vertex types [4].
2: Curved Objects.		
	a Agin	Description of several curved objects, such as Barbie doll and toy horse, by fitting generalized cylinders to range data [5].
	b Horn	Reconstruction of arbitrary curved surfaces, such as noses, from brightness profiles [6].
	c Turner	Interpretation of ideal line drawings as objects with second order curved surfaces. Interpretation of extracted line drawings as cylinders, cones and polyhedra [7].
II. INDOOR WORLDS.		
1: Faces.		
	a Kelly	Sequential location of prominent features e.g. eyes, nose, mouth, top-down in a predetermined contextually dependent order within crude face outlines extracted from a coarse image [8].

b Fischler	Simultaneous location of prominent facial features by finding optimal embedding of a flexible template [9].

2: Room Scenes.

a Garvey and Tenenbaum	Location of objects, e.g. door, tabletop, chair, in multi-sensory images using interactively and automatically generated distinguishing feature strategies[10].

III. OUTDOOR WORLD.

1: Landscapes.

a Lieberman	Location of prominent landscape features, e.g. sky, tree, water, top-down in a predetermined, contextually dependent sequence based on images crudely segmented into homogeneously colored regions [11].

2: Road Scenes.

a Yakimovsky	Segmentation of simple road scenes into regions yielding the optimal joint interpretation according to first order Bayesian semantics [12].

3: Cityscapes.

a Ohlander	Segmentation of outdoor scenes by recursively partitioning color and brightness distributions [13].

---

few constraints among objects, or even pieces of objects, limited to structural constraints, and, perhaps most significantly, the existence of unambiguous, goal-independent interpretations. The Indoor World, encompassing room scenes and faces, is characterized by functional names, strong contextual constraints, unambiguous interpretations and slightly richer sensory data, notably color. From these worlds it is a significant step to the Outdoor World of landscapes, in which occur overwhelming amounts of image detail, ambiguity, lack of contextual constraints, lack of functional names, variability of members of a given conceptual class and many possible varieties of context and goals for description.

It became apparent that, in order to produce a good description of a scene, arbitrarily complex reasoning may be necessary, involving many types and levels of knowledge. Considering the richness and complexity of the world, it was felt necessary to limit the range of difficulties to be faced by limiting the problem domain. This led to research into narrow systems with data about a very restricted class of scenes, for example, programs in the world of polyhedral objects know about gravity and support. By relying heavily on particular characteristics of the domain, such programs could afford to ignore knowledge and broad principles necessary for scene analysis in general. The selected characteristics are often caricatures of some limited aspect of reality, and are commonly implicitly embedded in the computer program in a way that makes if difficult to appreciate not only their significance for performance, but sometimes even that they are represented at all. Thus, scene analysis systems for the world of polyhedra are based on the straight line as an element of structure, while most landscape programs totally ignore all structure other than adjacency.

The performance and generality of programs within any domain depend on the amount of available knowledge about the domain, the amoung of flexibility available in using that knowledge and the degree of modularity with which new knowledge is incorporated. Winston [14] has documented this system in successive generations of blocks world programs written at MIT.

Roberts' line extraction program, for example, used implicit knowledge of the importance of straight edges in concatenating gradient points into line fragments, merging collinear fragments into lines and segmenting lines at corners. This knowledge and its rigid application in a sequence of independent local operations proved inadequate, some lines are missed, shadows, reflections and surface markings cause extra lines; errors made at one stage are compounded at later stages. These deficiencies motivated Shirai to invoke additional domain-dependent knowledge in the line finding process.

Shirai's program found edges, one at a time, at each step using heuristic knowledge about line drawings of polyhedral bodies implicit in the context of previous results to predict where additional lines would most likely be found. The program knew to look first for the outer boundaries of an object or cluster of objects, presuming high contrast between white blocks and black table, and then tried to guess lines that were extensions of, or were parallel or perpendicular to, lines already found. The search for predicted lines was concentrated around vertices where they were most likely to be found. The priority of search depended on the specificity of the prediction. The additional knowledge helped Shirai's program obtain significantly better performance and reliability.

Another interesting contrast is found in a comparison of the techniques for line drawing interpretation employed by Roberts and by Waltz. Interpretation to Roberts meant choosing a stored, three-dimensional model which, after translation, rotation and scaling, could be projected to match part of the image. The recognized piece was then cut away and the remainder considered. Unfortunately, many distinct decompositions of cluttered scenes and even of single objects could usually be obtained with a given set of models. Waltz's program avoided this difficulty by abandoning rigid models and relying instead on a large number of local descriptions of edges and vertices, plus rules for combining them, to form legal scenes of trihedral solids. (Analogies of this type are, of course, limited because the programs seldom share exactly the same objectives.)

Work of Kelly and Fischler on faces also provides some illuminating contrasts on how the use of knowledge affects program performance and generality. Both programs model faces with caricatures of prominent features, such as head outline, eyes, nose and mouth, whose relative positions are allowed to vary within prescribed limits. Each feature is represented procedurally in terms of its gross attributes (e.g. nose might be represented by a routine which tests for two dark circles corresponding to nostrils), whose meaning derives primarily from the global context. However, these programs use their common knowledge of faces in dramatically different ways. Kelly's strategy was to search the image top-down for a collection of edges that could correspond to a plausible head outline. This outline was used to predict approximate location in which to search for facial features using the crude procedural templates. The program looked first for the usually most reliable features, and used them to refine further the predicted locations for subsequent features. Global confidence accrued as additional features were verified. The eyes were sought first in an area defined with respect to the top and sides of the head outline. If this search succeeded, the nose was sought next in an area predicted with respect to the eyes. Then the mouth was sought in an area delimited vertically by the nostrils and bottom of the head outline, and horizontally by the two eye centers.

A major shortcoming of Kelly's and many similar programs, was that features were sought in a predetermined, contextually dependent sequence. Hence a single obscured feature could derail the entire analysis. A second, and related, drawback was lack of generality. Knowledge about objects was deeply embedded in procedures, often in the form of ad hoc heuristics and thresholds. Developing these procedures was a tedious empirical process that must be repeated for another narrow domain. Furthermore, even small modifications can involve extensive reprogramming; a single new object might substantially alter the features needed to distinguish previously defined objects as well as the contextual order in which objects should be acquired.

Fischler avoided the above difficulties by abandoning the notion of sequential search and evaluating simultaneously the best overall embedding of a face prototype in the image. In Fischler's program, the facial model was used as data for a general optimization procedure which treated it as an elastic template. Unlike Kelly, the same optimization procedure could be used with different data to interpret a variety of domains.

Lieberman's program for interpreting landscape scenes used a sequential, context-driven strategy, similar to Kelly's, while Yakimovsky's program used a global optimization approach similar to Fischler's.

We mention, in passing, that another way of overcoming the above limitations is to develop a system that can plan its own strategies based on global knowledge about a particular domain and specific knowledge about the current scene. Such a system could alter its strategy dynamically to capitalize on the strongest available evidence. Moreover, it would have the inherent generality of being able to function in any environment for which its knowledge base and perceptual primitives are adequate. Garvey's work in room scenes is a first attempt at automating the development of perceptual strategies in this way.

More important than particular programs, research to data has given us some understanding of some of the design principles for visual systems and an appreciation of the type of knowledge they must possess in order to function well.

What knowledge is needed?

Relevant knowledge can have only a positive effect upon the competence of a scene analysis system. The more the better, apparently without limit. Adding (correct) information enables interpretation in some cases in which it was not previously possible, or corrects some erroneous analyses; deleting information has the reverse effect. There appears to be no firm boundary between seeing and reasoning, in order to do the former, various amounts of the latter may be necessary.

This precept finds support in the manifest disadvantages of a parsimonious description, an object may be unrecognizable if a single crucial feature is obscured. It is generally both more economical and more reliable to characterize objects redundantly by a juxtaposition of many, perhaps crude, features than by detailed but irredundant description of a single feature. Moreover, such a juxtaposition may provide the only feasible description for highly variable objects. A tree, for example, is more naturally characterized by the gross attributes "large", "green" and "irregularly shaped" than by a detailed description of any single attribute. The desirability of redundant data in the form of multiple or multi-sensory images follows as a consequence of the desirability of redundant description [15].

A competent scene analysis system must embody, explicitly or implicitly, knowledge about the process that relates the picture to the scene. An understanding of this process involves knowledge of projective geometry, topology, optics and electro-optical transducers, as well as information about surfaces and lighting, including reflectivity, specularity, highlights, color, surface texture, shading etc. This concern with the physical basis of the picture is one of the principal distinctions between Scene Analysis and pictorial Pattern Recognition.

The above knowledge explains the relation between a patch of surface and the corresponding patch of picture, but patches of surface are not arbitrarily arrayed: surfaces may be continuous, or not, and may occlude one another. Structure is important and manifests itself in terms of objects and their features, bumps and hollows – their shape. We normally describe the world in terms of objects, and we wish a scene analysis system to do the same. Note that "object" is a fundamental notion, imposed by the observer, and is context-dependent. Often, the evidence for attributing two regions of a picture to different objects may be weaker than evidence for further subdividing one of the regions. Only experience and context allow the right decision. For example, consider a black and white cat sleeping on a black fur rug. Thus the system must know about likely objects and parts of objects. It should have generic knowledge of cats, as well as of specific cats, in order to deal with new cases of the same sort.

Once the concept of *object* has been introduced, it follows that knowledge of relationships between objects, the ways in which they interact, the ways in which they behave, likely events and so on, all helps in the analysis of pictures. For example, *gravity* and *support* help in judging relative sizes. Since two men cannot be floating in the air, their feet must be resting on the field, which means that man whose feet are higher in the picture is further away. The fact that he also appears larger must be reconciled by concluding that the other man is short. Naturally, we are not suggesting that such reasoning actually occurs when we see, at least not in this form - rather, that knowledge often has ramifications which result in the system behaving as though such reasoning had occurred.

Reddy [16] has listed some examples of knowledge that can be employed in scene analysis, and has categorized them in an interesting way. A scene can be described hierarchically at the *feature* level as a vector for each *pixel*, e.g. intensity, hue, saturation, depth, etc., at the *segment* level as a contiguous grouping of pixels which do not differ significantly in one or more feature dimensions, at the *regions* level as sets of segments grouped over shadows, occlusions and highlights, at the *object* level as several regions belonging to one or more objects, at the *structural* level as ensembles of objects which obey known relationships (e.g. table and chairs), and finally at the *environment* level as a single, meaningful description of the whole scene. In addition to specific knowledge mentioned by Reddy, there is also a substantial amount of metaknowledge concerned with applicability and reliability in particular contexts. For example, the rule "sky is blue" depends on implicit assumptions of time of day, season, weather, geography, location and so forth.

Reddy views knowledge as a set of rewriting rules which transform an image represented at one level into a corresponding representation at another, higher level. Two examples of knowledge expressed in

Reddy's hierarchy are "sky is blue" (structure/feature) and "city scenes involve vertically oriented rectangular shapes" (environment/segment). Note that knowledge may transcend several levels, and the logical structure of the scene does not dictate the processing sequence. A good example is the acknowledged interdependence of segmentation and interpretation, since the significance of pictorial entities can only be evaluated in the broader context of interpretation. In fact, it no longer seems appropriate to view segmentation and interpretation as distinct processes.

How are knowledge and descriptions represented?

A number of issues and non-issues have arisen over the representation of knowledge, not just visual knowledge. It seems desirable to have a redunant representation because data organized in a single way may be accessed easily for some queries but very awkwardly for others. To a large extent, however, this is a non-issue, a detail of implementation rather than fundamental science. More to the point is the ability to manipulate the representation easily, extract implicit information as required and the necessity for cross-links so that relevant related information is readily available when needed.

The non-issue of Procedural Versus Assertional has been aired in other domains than vision [17-18]. We simply remark here that one mans procedure is another mans data. An interpreter, a theorem prover or a (CPU) Control Processing Unit are all operating on data (assertions), or all running procedures. What is at stake are the practical issues of space and time. To the seeker of truth they are only important insofar as they affect his search. Compiling (procedural embedding) the knowledge makes running faster, at the expense of making modification slower. It seems that natural vision systems may involve some compilation of frequently used processes, as well as retaining the fundamental ability to work from scratch using a reasoning ability, unfamiliar objects take some time to perceive until they become familiar.

A similar issue is Symbolic versus Iconic. Should the information be encoded by symbols that stand for generalized stereotypes, or should it stand for itself? Since it is only by abstracting from the current particlar case that general principles may be invoked, it is clear that symbols are vital. However, they lose touch with details of reality, so the original data must be maintained in some form to handle questions whose answers are not embodied in the original abstraction. The relation between symbols and images is very much an open question. It may be that one should convert to symbolic representation at a very early stage [Marr 19]. The evidence concerning human representation is mixed; much of our visual imagery is in terms of stereotypes, rather than the original images, as though we reconstruct images from symbolic descriptions and associated iconic examples [Norman 20]. However we are quite good at remembering details, which hints against highly abstracted symbolic descriptions.

Elementary considerations of flexibility in the use and acquisition of knowledge dictate that, whatever the form of knowledge, be it procedural, assertional or iconic, it should be made explicit, rather than buried in the coe and forgotten.

How should the knowledge be used?

The use of knowledge in vision is determined by the following observations:

Because of noisy data and probabilistic constraints, vision is not a purely deductive problem with a unique, "correct" solution. It requires determining a "best" interpretation of the data in the light of the available knowledge. The situation is exacerbated by the quantities of data and knowledge involved, by the fact that the knowledge sources are error-prone and not unanimous in their recommendations, and by the strong interactions between them. (This interaction is manifest, for example, in the ways in which segmentation and interpretation are inextricably intertwined.) These considerations preclude a straightforward search for the best solution; ways must be found to circumvent these problems.

Naive backtracking, i.e. returning to a previous choice point and taking a different alternative, is utterly hopeless. It appears vital to take advantage of local independence of parts of the picture and only undo the direct consequences of an erroneous choice. Perhaps vision has something important to contribute to problem-solving here [21]. It is even better to avoid redundant work in the first place by eliminating possibilities as soon as possible rather than eliminating them many times over on different branches of the search tree. Waltz's filtering algorithm provides a dramatic illustration of this point. A generalization of this principle for optimization problems is to make locally-best choices which, if based on enough information, can be expected to lead to a near optimal solution with minimal search.

The assessment of local evidence requires integrating many sources of knowledge, of different types and levels, which very often yield contradictory evidence, e.g. two adjacent patches of a picture may have the same color, which suggests they belong to the same object, but may have different brightnesses, which suggests they do not. Higher-level knowledge about shadows indicates that the patches are from an object with a shadow across it, but perhaps knowledge of illumination indicates that the shadow could not be cast in such a position unless light had been reflected from some surface, and so forth.

It would appear that some sort of consensus mechanism is necessary, an environment in which knowledge sources can compete and cooperate to yield a result upon which they can largely agree. The relative weightings attached to evidence from different sources of knowledge depend dynamically on the current processing context and ultimately on the data, the expectations and the goals.

Some of these principles are illustrated in a working program, MSYS, recently developed at SRI, for assigning interpretations to regions in previously segmented office scenes [22]. The problem here was to find the set of region interpretations having the highest joint likelihood. Each possible interpretation is represented by a process which computes its global likelihood based on local region attributes, and the likelihoods of other contextually related interpretation. These processes interact, compete and cooperate, adjusting likelihoods until they achieve a state of dynamic equilibrium that corresponds to the best solution consistent with the evidence. The equilibrium state is, by its very nature, reversible and processes can be defined so that information is never lost, but can be recomputed if the consensus demands it. Because of the manner in which processes communicate, via a global repository of data, it is possible to add or remove them (dynamically if desired) with incremental changes in system performance. Hence the system can readily be modified or extended in a modular way.

The use of an equilibrium model eliminates some issues of sequencing, but does not completely sidestep the need to strike a balance between data-driven and goal-driven organizations. The relative importance of observations and expectations still depends on goals. Data-driving has the advantage that it is less likely to get out of touch with reality but also has the disadvantage of considering too many possibilities. Goal-driving considers only relevant possibilities, at the expense of ignoring the unexpected. Evidently some combination is desirable, such as an hypothesis-verify paradigm, but the best organization is still in doubt.

The organization issue also arises in representing relevant knowledge and hypotheses within the network of equilibrium processes. As a practical matter it must be done dynamically as relevance emerges.

How should knowledge be acquired?

The substantial amount of World knowledge required for perception is most reasonably acquired in an incremental fashion as deficiencies are discovered, instead of trying to do an exhaustive case-analysis beforehand. Certainly as a matter of experimental convenience, incremental improvement ranks high. Humans certainly increment their visual skills, learning finer discriminations or broader generalizations as circumstances warrant.

Whether knowledge is acquired by supervised or unsupervised learning is not at issue here. As a practical matter, in the near term, it will be important to have good interactive facilities for empirically determining perceptual concepts and processes and communicating them to a machine [Tenenbaum 23].

What have we learned?

Perhaps the most important thing we have learned from our experiences in vision research is that it is a difficult and complicated task. Hopefully, we have also learned why it is hard. All the key AI issues crop up – world models, representation, organization, reasoning, problem solving, even manipulation. The full range of reasoning, from simple association to complex deductions, with uncertain evidence and uncertain rules is involved. The organizational problems involved with handling the large amounts of data and knowledge that arise with real pictures rule out Toy solutions.

The limited abilities of current Scene Analysis systems can be traced to compromises necessitated by the magnitude of the vision task. The quantities of data involved are huge but are still much less than is desirable. Many problems encountered in processing $64 \times 64 \times 4$ bit pictures are greatly eased (though not necessarily eliminated) by increasing resolution and number of brightness levels, particularly if the digitization is performed with something other than a noisy TV camera.

The knowledge possessed by existing systems is very limited. For example, there are several Blocks-World systems that specialize in some details of the problem; Waltz's catalogue, Shirai's block-finder etc., but there is no unified system that contains and uses all we currently know about the Blocks domain.

The knowledge of systems is frequently and rigidly embedded in procedures which use it in a limited way. For example, knowledge that is used passively, as a recognition test, may not be available for use actively in searching for the object [24].

What should be done?

Existing low-level abilities seem over-specialized to particular domains. Their usual dependence upon one characteristic, to the exclusion of others leads to marginal and brittle performance, weakness in forming descriptions of novel situations (which hampers learning), and the need to tear apart and rebuild for each new domain. Work in scene analysis is beginning to address some of the issues raised in this article. There is an important need for a better foundation on which to build the higher functions of scene analysis systems that are experts in their chosen domains.

The time is right to develop broadly-based scene analysis systems. By this we mean systems that have a great deal of low level knowledge, about surfaces, reflectivity, shading, color, illumination, occlusion, shadows and so forth, that is common to many domains. Such systems would be able to make a good attempt at segmenting and interpreting (in simple terms) an arbitrary scene. Human beings can almost always make some sense of a scene, even when it comes from a domain with which they are unfamiliar (for example, photomicrographs or abstract art). They can at least identify significant lines and regions even when the evidence for them is quite weak, and despite their apparent lack of meaning. No existing machine can approach human performance in such situations.

An unified framework is needed for experimenting with both general and domain-specific knowledge in scene interpretation, where the quality of the result can be progressively improved by incorporating additional knowledge. With just low-level knowledge, the system should manifest behavior consistent with the Gestaet laws of human vision. Augmented with domain-specific knowledge, the framework should rival the competence of the best current special-purpose systems and be more readily extensible and considerably more robust.

The proper organization of such a broadly-based system involving many conflicting sources of knowledge is a key issue to which we have already devoted some attention above. This issue has also received considerable attention in the literature of *Scene Analysis* [14][25] as well as in other areas of AI, notably speech understanding (ref. HEARSAY II [26]). A consensus has begun to form on some specific design criteria, such systems should be data driven; all knowledge should be modularly represented in a shared global data base, and competing hypotheses should be explicitly represented and freely available in the data base, rather than hidden in the internal backtracking variables.

This type of organization, as exemplified by MSYS and HEARSAY, has so far only been used with domains – specific knowledge for interpretation of segmented scenes. We believe it could be used with more general knowledge to assign consistent low level interpretations (e.g. of color, relative depth, orientation) to edge fragments and small sets of picture elements and thereby integrate segmentation with interpretation.

In addition to the specific suggestions proposed above, we offer a few general guidelines based on lessons of the past decade:

1) Researchers concerned with the scientific issues of vision in general should acknowledge that the prodigious quantities of data and knowledge involved preclude any hope of near real-time operation on current

serial computers. They should therefore avoid the preoccupation with efficiency that characterized the ingenious but special purpose solutions described earlier.

2) A corollary of (1) is always to work with the highest quality data that can be obtained. In the past, many researchers have settled for degraded, low resolution images which fit conveniently in core, and rationalized their acceptance with arguments that this data could be correctly perceived by humans. However, the problem is probably being made unnecessarily difficult by eliminating local detail (which might allow local, rather than global interpretation). Humans can cope with such bad pictures but perhaps it is premature for machines to do so.

3) All assumptions used in the research should be carefully documented so that the work can be meaningfully evaluated, both with respect to its own objectives, and in comparison to related work with similar objectives. Examples of assumptions would include the quality of data, the types and quantities of knowledge required, processor requirements, the overall objective (general purpose or special system), and detailed assumptions about the types of objects and events in the scene domain. Libraries of standardized scenes should facilitate realistic comparisons of alternative techniques.

4) Work should proceed simultaneously in several pictorial domains to force generality and provide opportunities for cross-fertilization.

5) Motion perception has been seriously neglected, despite its paramount importance in human perception. Sequences of pictures should be processed to explore the constraint of continuity.

6) Systems should be designed to assimilate advice or new knowledge provided incrementally by the experimenter to prevent recurrence of errors. Hopefull this input can be provided via pictorial examples. It is more natural to designate pictorial concepts by pointing with a cursor than by editing a program.

Conclusion

We close by reiterating some of the major premises underlying this paper:

*** The more knowledge, the better.
*** The more data, the better.
*** Vision is a gigantic optimization problem.
*** Segmentation is low-level interpretation using general knowledge.
*** Knowledge is incremental.
*** Research should pursue Truth, not Efficiency.

Another decade will determine our skill as Visionaries.

References

[1] L. G. Roberts, "Machine Perception of Three-Dimensional Solids," MIT Lincoln Lab., Lexington, MA., Tech. Report No. 315, May 1963. Also in *Optical and Electro-Optical Information Processing*, J. T. Tippett et al., eds., MIT Press, Cambridge, MA., 1965.

[2] G. Falk, "Interpretation of Imperfect Line Data as a Three Dimensional Scene," *Artificial Intelligence*, Vol. 3, No. 2, 101-144, 1972.

[3] Y. Shirai, "A Context Sensitive Line Finder for Recognition of Polyhedra," *Artificial Intelligence*, Vol. 4, No. 2, 95-119, Summer 1973.

[4] D. G. Waltz, "Generating Semantic Descriptions from Drawings of Scenes with Shadows," AI TR-271, Artificial Intelligence Laboratory, MIT, Aug. 1972.

[5] G. J. Agin and T. O. Binford, "Computer Description of Curved Objects," *Adv. Papers 3d Intl. Conf. on Artificial Intelligence*, Stanford Univ., Stanford, CA., Aug. 1973.

[6] B. K. P. Horn, "Shape form Shading: a Method for Obtaining the Shape of a Smooth Opaque Object from One View," MAC Tech. Report 79, Project MAC, MIT, Cambridge, MA., Nov. 1970.

[7] K. J. Turner, "Computer Perception of Curved Objects Using a Television Camera," (Ph. D. thesis), School of Artificial Intelligence, Edinburgh University, Edinburgh, Nov. 1974.

[8] M. Kelly, "Visual Identification of People by Computer," Memo. AI-130, Comp. Sci. Dept., Stanford University, Stanford, CA., July 1970.

[9] M. A. Fischler and R. A. Elschlager, "The Representation and Matching of Pictorial Structures," *IEEE Trans. on Computers*, Vol. C-22, No. 1, 67-92, Jan. 1973.

[10] T. D. Garvey and J. M. Tenenbaum, "On the Automatic Generation of Programs for Locating Objects in Office Scenes," *Proc. 2d Intl. Joint Conf. Pattern Recognition*, 162-168, Aug. 1974.

[11] R. Bajcsy and L. I. Lieberman, "Computer Description of Real Outdoor Scenes," *Proc. 2d Intl. Joint Conf. Pattern Recognition*, 174-179, Aug. 1974.

[12] Y. Yakimovsky and J. A. Feldman, "A Semantics-Based Decision Theory Region Analyzer," *Proc. 3d Intl. Conf. on Artificial Intelligence*, Stanford University, Stanford, CA., Aug. 1973.

[13] R. Ohlander, CMU Thesis.

[14] P. H. Winston, "The MIT Robot," *Machine Intelligence*, Vol. 7, 431-463, B. Meltzer and D. Michie, eds., Edinburgh University Press, Edinburgh and American Elsevier Publishing Company, 1972.

[15] J. M. Tenenbaum, "On Locating Objects by their Distinguishing Features in Multi-Sensory Images," Computer Graphics and Image Processing, Dec. 1973.

[16] A. Newell and D. R. Reddy, "Image Understanding: Some Notes," Minutes ARPA Image Understanding Workshop, Science Applications Inc., March 1975.

[17] T. Winograd, "Five Lectures on Artificial Intelligence," AIM-246, Stanford Univ. Artificial Intelligence Laboratory Stanford, CA., Sept. 1974.

[18] G. J. Sussman, "A Computational Model of Skill Acquisition," Tech. Note AI TR-297, Artificial Intelligence Laboratory, MIT, Cambridge, MA., Aug. 1973.

[19] D. Marr, "On the Purpose of Low-Level Vision," AI Memo 324, Artificial Intelligence Laboratory, MIT, Cambridge, MA., Dec. 1974.

[20] D. A. Norman and D. E. Rumelhart, *Explorations in Cognition*, Freeman and Company, San Francisco, April 1975.

[21] M. Minsky, "A Framework for Representing Knowledge," in *The Psychology of Computer Vision*, P. H. Winston, ed., McGraw-Hill Book Company, New York, 1975.

[22] J. M. Tenenbaum and H. G. Barrow, "MSYS: A System for Reasoning about Scenes," forthcoming SRI Artificial Intelligence Center Tech. Report, Stanford Research Institute, Menlo Park, CA.
[23] J. M. Tenenbaum, et al., "An Interactive Facility for Scene Analysis Research," SRI Artificial Intelligence Center Tech. Note 84, Stanford Research Institute, Menlo Park, CA., Sept. 1973.
[24] E. C. Freuder, "Active Knowledge," Vision Flash 53, Artificial Intelligence Laboratory, MIT, Cambridge, MA., Oct. 1973.
[25] S. E. Fahlman, "A Hypothesis-Frame System for Recognition Problems," Working Paper 57, Artificial Intelligence Laboratory, MIT, Cambridge, MA., Dec. 1973.
[26] L. Erman and V. Lesser, "A Multi-Level Organization for Problem Solving Using Many Diverse, Cooperating Sources of Knowledge," Dept. Comp. Sci., Carnegie-Mellon Univ., Pittsburgh, PA., March 1975.

(Continued from Page 2)

About 30,000 gallons of No. 6 oil, similar to Bunker C, was poured into one hold of the barge. This hold and a hold across the deck from it were insulated and fitted with heat exchange plumbing, providing temperature control capabilities. The oil at various temperatures was transferred by a submerged pump through a piping system on deck above the oil compartment into the empty compartment. After reaching a new experimental oil temperature the oil was pumped back into the original hold. About 15 to 20,000 gallons of oil was pumped during each of the 24 transfers to date.

Two pumping units have already been tested and evaluated, according to Mittleman, and another is presently undergoing evaluation. One of these units was able to pump 150 gallons of oil per minute at low temperatures and 1100 gallons per minute at high temperatures. Mittleman is being assisted in the program by Mark Jones.

The program is continuing and NCSL now has the only facility of its kind that has the capability to pump oil under controlled conditions—controlling temperature, measuring the flow rate and measuring the pump pressure.

The program being conducted here is under the sponsorship of the Supervisor of Salvage. The experience produced from this effort will provide the Navy with the basic data for supporting salvage decisions involving oil cargoes potentially hazardous to the environment.

El Nino Expedition — A Step Toward Ocean Prediction

William C. Patzert*
Scripps Institution of Oceanography

For the past few years, a group of scientists, sponsored jointly by the Office of Naval Research and the National Science Foundation, has been studying large-scale (> 1,000 km), long-term (months to years) ocean-atmosphere fluctuations with an ultimate objective of improving prediction of climate and weather for the Pacific Ocean and North America. This project, the North Pacific Experiment (NORPAX), is presently conducted at 18 different institutions by 34 scientists. Recently, a group of NORPAX scientists conducted a five-month expedition to verify their prediction of an unusual event in the ocean and atmosphere.

The event, El Nino, is characterized by unusual temperature changes over large areas of the eastern tropical Pacific Ocean. These large-scale temperature anomalies in the upper ocean have been shown to be manifestations of changes in the ocean-atmosphere system over the entire Pacific Ocean (Bjerknes, 1961) and possibly over the entire globe (Bjerknes, 1974). The unusual warming, referred to as El Nino, has dramatic consequences both on the coastal and offshore environs of Ecuador and Peru. Offshore, the normally cold, upwelling coastal waters which join the cool, northward-flowing Peru-Chile current are invaded by much warmer waters from north of the equator and the west. These warm waters induce tropical-like rains that cause frequent coastal flooding and, offshore, disperse the Peruvian Anchovies, normally the world's largest single fishery. In all, El Nino is a natural catastrophe for Peru.

For NORPAX scientists, the study of El Nino has provided one of the focal points for probing the larger-scale causes of climate. Based on our understanding of previous El Nino events, a prediction was made in October 1974 (Quinn, 1973) that a weak El Nino event would occur in early 1975. This prediction was based on the southern oscillation index, which is a long time-series of the difference in sea level pressure between Easter Island, representing the South Pacific sub-tropical high, and Darwin, Australia, representing the Indonesian equatorial

Dr. Patzert is a Research Associate at Scripps Institute of Oceanography, University of California, San Diego, specializing in descriptive physical oceanography.

low. Our group, led by Dr. Klaus Wyrtki of the University of Hawaii, quickly organized a multi-disciplinary expedition to observe conditions offshore from Ecuador and Peru. The proposal was submitted in November and the decision to proceed was made in December. Fortunately, both the Office of Naval Research and the National Science Foundation overcame bureaucratic hurdles and made funds and equipment available on short notice.

In order to describe the Nino occurrence and its development with time, the area (2°N to 14°S, 95°W to the coast) was surveyed twice (mid-February through March and mid-April through May) (Figure 1a and 1b). The observation program on board the University of Hawaii research vessel MOANA WAVE (Figure 2) consisted of vertical profiles of temperature and salinity to 500 m depth at stations 60 nautical miles apart, intermediate temperature profiles by expendable bathythermograph XBT at 20-mile intervals, and hydrographic stations with bottle casts for oxygen and nutrients at 120-mile intervals. In addition, zooplankton and phytoplankton samples were obtained, and primary productivity and chlorophyll-A were measured. Twice-daily upper air observations with radiosondes were collected by Navy meteorolgists and surface meteorological observations at 2-hour intervals were recorded. Personnel aboard the R/V MOANA WAVE represented many institutions: Scripps Institution of Oceanography, University of Hawaii, Duke University, University of Washington, Instituto del Mar del Peru, Instituto Oceanografico de Ecuador, and the Office of Naval Research.

During the first cruise, Southeast Trade winds were extremely weak north of 10°S. In the ocean, a transgression of warm, low-salinity water across the equator to 5S was observed (Figure 1a). Although weak upwelling was still present along the coast of Northern Peru, no cool water was advected northwest from this upwelling region, as in normal years. The warm water flowing southward between the Galapagos Islands and Ecuador was confined to a shallow surface layer 10 to 25 m thick. The changes in the subsurface layers were also quite different from what is normally expected. The density surfaces along both the equator and the coasts of Ecuador and Peru were abnormally depressed in the upper 200 m, indicating a very strong eastward flow of the equatorial undercurrent and southward flow along the coast. This anomalous flow regime seems to play an important role in El Nino dynamics.

While observations made during the first cruise revealed features of a developing El Nino, the data of the second cruise documented that it was short-lived. During April-May, Trade Winds quickly returned, and the cooler ocean temperatures normally expected were quickly returning. Upwelling along the coast and the equator was re-established, and the cool waters of the Peru-Chile Current had replaced the warm

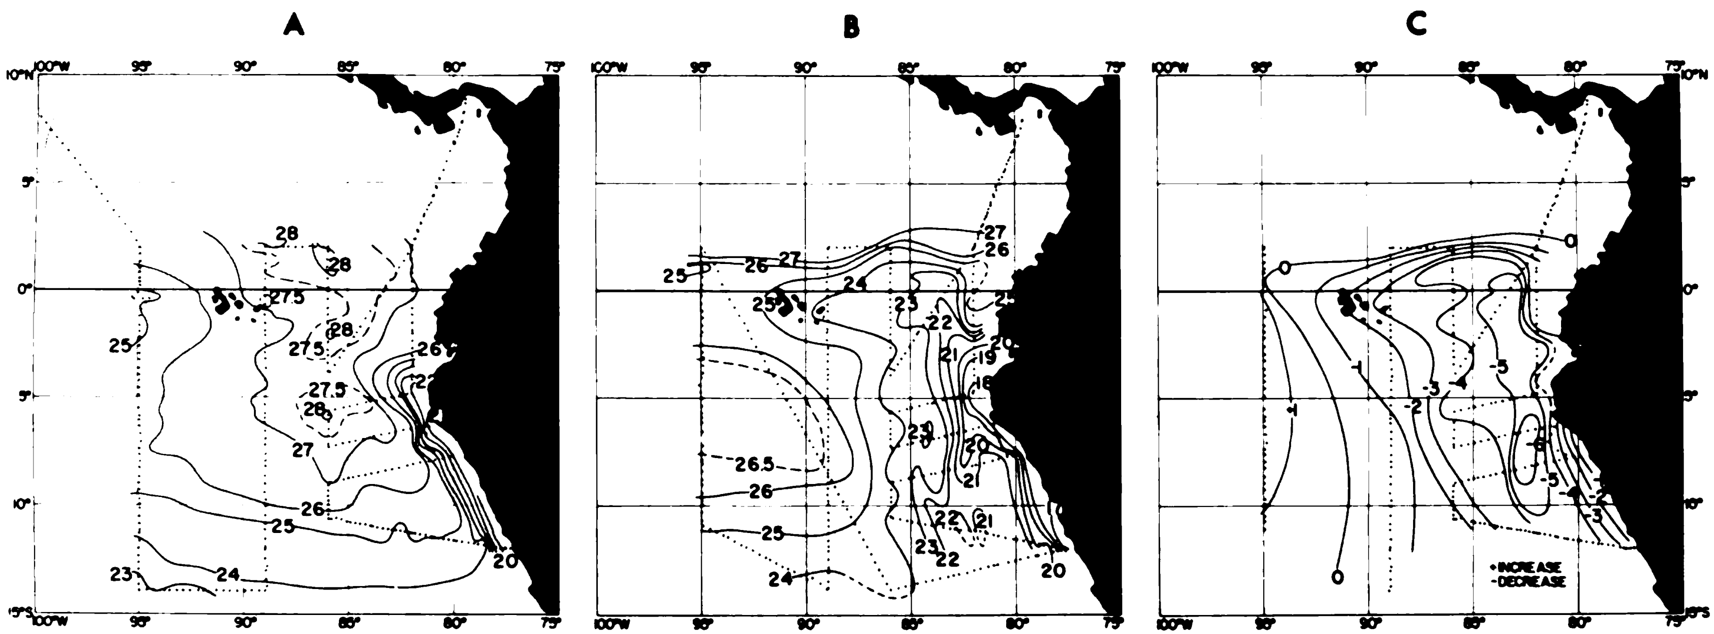

Figure 1 — a) Sea surface temperature (°C) from 11 February to 31 March 1975, b) sea surface temperature (°C) from 17 April to 27 May 1975 and c) changes in sea surface temperature (°C) between the two El Nino cruises.

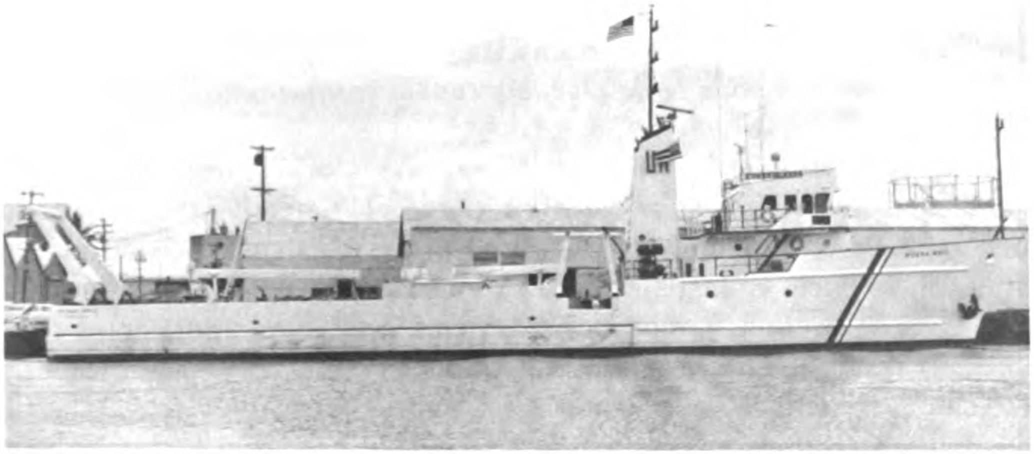

Figure 2 — The 174-foot R/V MOANA WAVE (AGOR-22) of the University of Hawaii is the latest in a series of research vessels built by the U.S. Navy for the academic research fleet. Powered by two diesel engines, the WAVE has a range of 8,000 miles and is able to stay at sea for 40 days.

waters found two months earlier (Figure 1b). Beneath the surface, the eastward flow along the equator had weakened and, at the coast, the flow was now to the north. The large change in the sea surface temperature between the two cruises (Figure 1c) illustrates how rapidly and intensely the ocean changes in this region. The change in sea surface temperature within this two-month period (greater than 5°C in some areas) was as large as the mean seasonal range indicated by climatological atlases.

Although 1975 was not the year of a catastrophic El Nino of long duration, the prediction was successful in that an El Nino event continued to develop from January through March. Previous to these cruises there had never been an expedition covering the entire area affected. Consequently, we have learned a great deal about the dynamics of El Nino events. The large ocean temperature and circulation anomalies observed and the rapidity of their dissipation have provided material for much careful thought. In the coming months, we will be analyzing these cruise data in concert with the larger-scale atmospheric and oceanic events that preceded this small El Nino in order to better understand the dynamics of Pacific climate.

Mid-Ocean Dynamics Experiment (MODE)*

Kathy Daly†
Woods Hole Oceanographic Institution

Scientists have long recognized that the world's oceans play an important role in global weather patterns. Ultimately, they want to understand the coupling of processes in the ocean with those in the atmosphere well enough to predict long-range weather and climatic changes. At present, atmospheric systems are better understood than oceanic systems. The large patterns of ocean circulation—for example, the clockwise circulation of water in the upper part of the North Atlantic are fairly well understood, but relatively little had been learned about the more local oceanic phenomena until the Mid-Ocean Dynamics Experiment (MODE) provided a description of one of these local events.

MODE focused on surprisingly energetic, rotating water colums which had been known for some years to exist throughout the world's oceans but for which oceanographers had no satisfying explanation. Sometimes thought of in terms of the weather systems of the deep sea, these ubiquitous motions—called "eddies"—are somewhat similar to the atmospheric high- and low-pressure systems which move within, yet strongly affect, world weather patterns.

In 1970, a group of scientists gathered together by Henry Stommel, of the Massachusetts Institute of Technology (MIT) and Woods Hole, initiated plans to "catch" an eddy and describe it in detail. Their plan became one of the largest physical oceanographic experiments ever undertaken; three years later, an international team of 50 oceanographers, representing 15 oceanographic laboratories and institutions, combined forces to find and document the life and characteristics of an eddy. Employing six ships, three airplanes, and extensive instrumentation in a 600-kilometer-wide area situated between Bermuda and Florida, MODE scientists were successful in mapping and describing not only one but a variety of eddies (Figure 1).

There is no simple physical definition for an eddy. While it moves through the ocean, it may change shape, intensifying or dissipating, gathering energy and releasing it. However, MODE scientists discovered

*This article is condensed from Woods Hole Notes, May 1975.

†Miss Daly is a student at the University of Massachusetts at Amherst. She was the co-editor of MODE Hot Line News when working at Woods Hole.

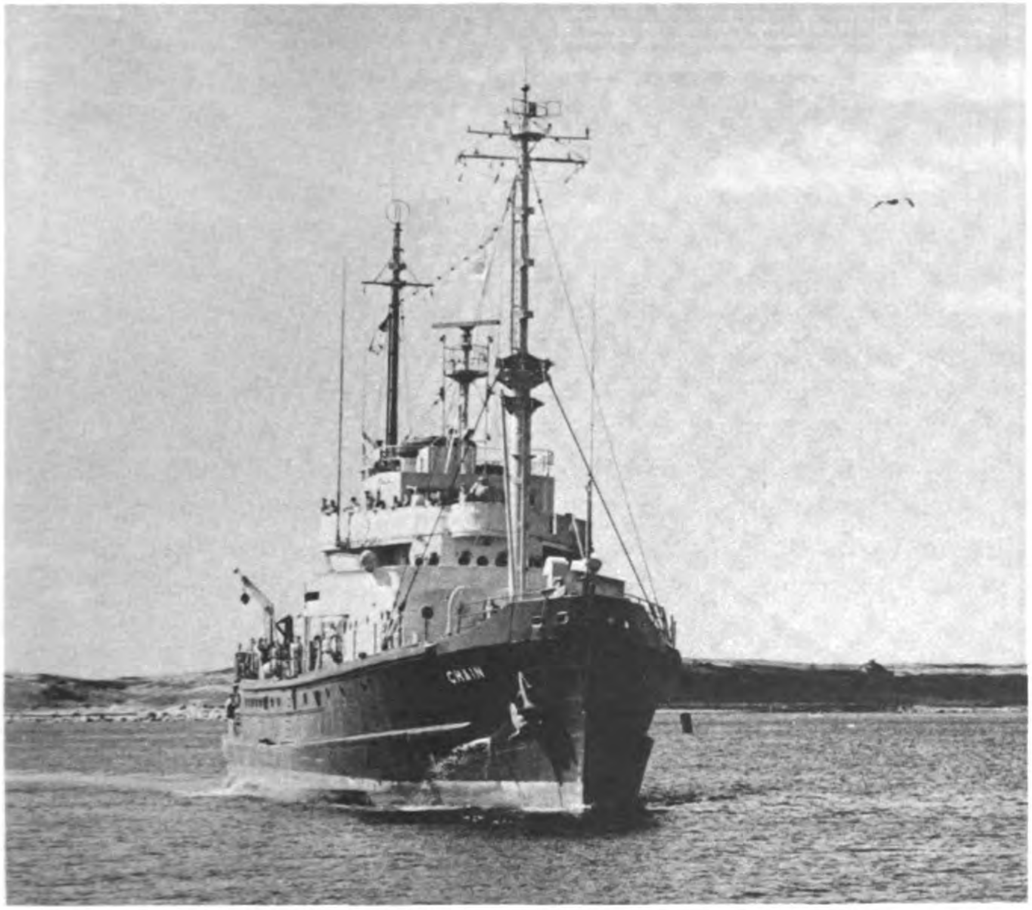

Figure 1 — R. V. CHAIN one of the six ships participating in the MODE field program

at least two distinct forms of eddies: slowly varying, but intense upper-layer (to 1,500 meters) eddies; and quickly changing, but less energetic, deep-water eddies. Peter Rhines, a Woods Hole Oceanographic Institution (WHOI) Senior Scientist and one of 15 Institution-staff MODE participants, said, "This confirms the fact that the behavior of the deep ocean has its own distinct properties and is not a mirror of what occurs in the shallower levels of the ocean."

Understanding and mapping eddy motions is complicated by their variability, which was more severe than anticipated. "The energy of the deep-water eddies, for example, changed by a factor of two in 200 kilometers, and we could see this feature persisting over two years' observations," Rhines says. In the MODE area, eddies were tightly bunched and generally moved westward through the area (Figure 2). However, a combination of mechanisms, such as variations in seabottom topography and boundary features (like the Gulf Stream), can contribute to a range of eddy characteristics. Thus, Rhines concludes, "What we saw of eddy motions in the MODE area may not be representative of other eddies in the Atlantic Ocean basin" (Figure 3).

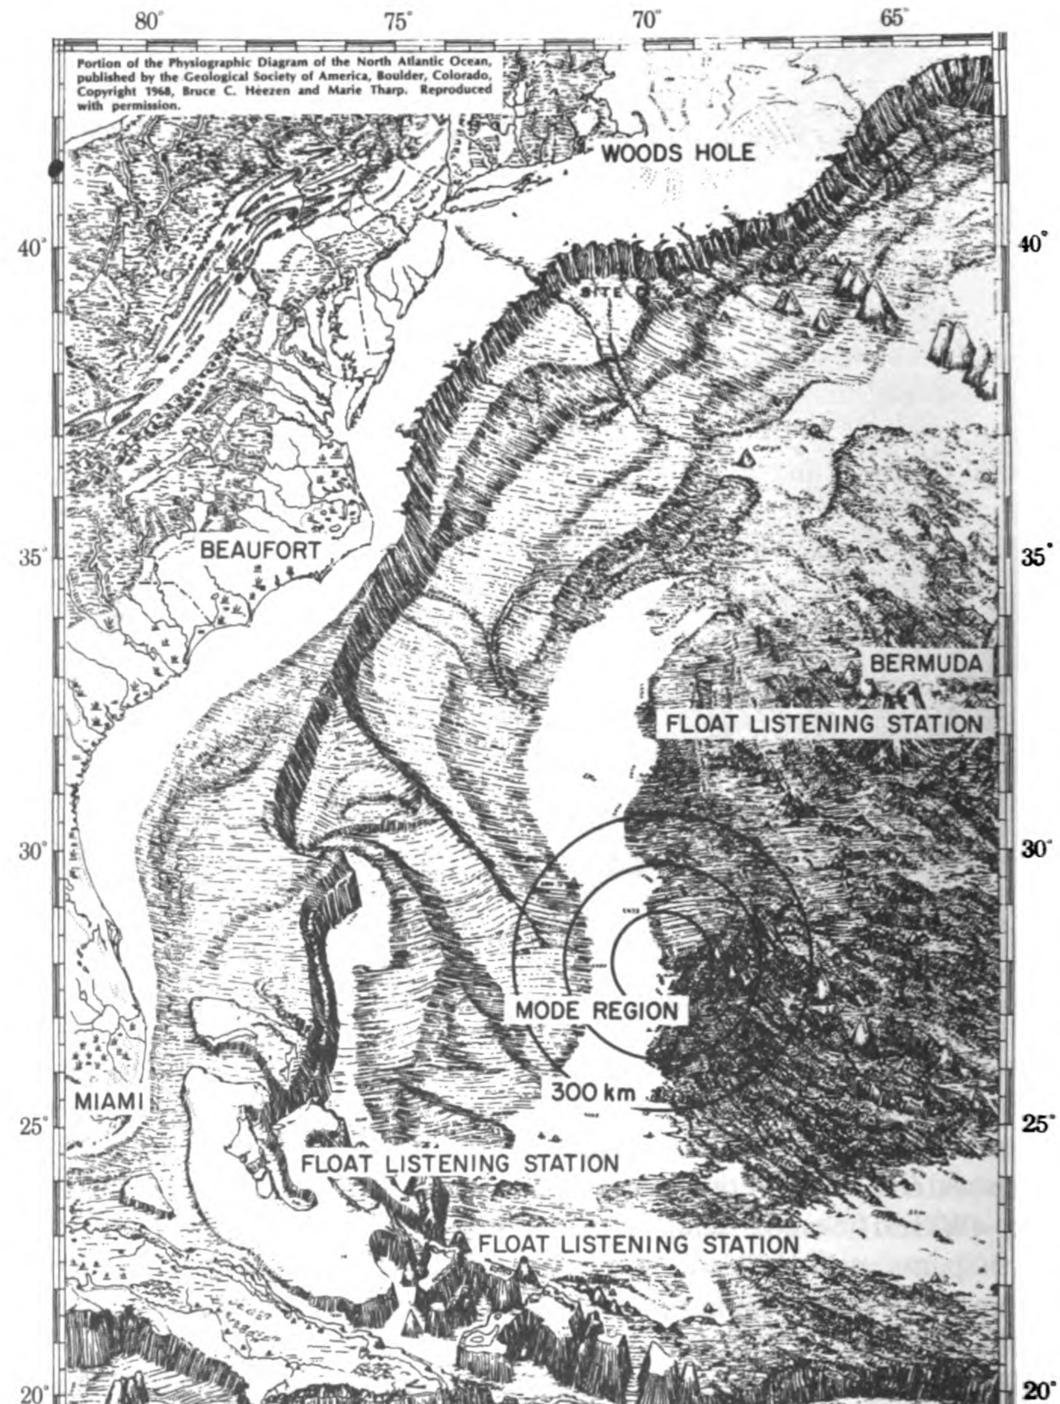

Figure 2 — Rings on map, a portion of the Heezen and Tharp physiographic diagram of the North Atlantic (1968), show location and topographic features of the MODE area. The moored instrumentation lay within the inner and middle rings (100 and 200 mile radii), and hydrographic stations were made out to the largest ring (300 mile radius). The POLYMODE studies will center around 55 degrees west and extend from 10 to 40 degrees north.

Figure 3 – Launch of towfish during MODE Project aboard CHAIN in 1973. Instrument was used to map density surfaces w/fine horizontal resolution.

First Indications

The first indication of the existence of eddies came from British Oceanographer John Swallow's expedition with the Oceanographic Institution's ketch ARIES in 1959 to measure the slow, broad drift of the deep waters in the western North Atlantic. He launched free-drifting floats, now known as Swallow floats, ballasting them to remain at 2,000- and 4,000-meter depths, and tracked their movement. Instead of finding, as he expected, slow-moving deep currents, he found currents moving with irregular, eddy-like motions and with speeds 10 times greater than he had anticipated.

Technical limitations prevented further exploration of this surprising phenomenon at the time, but the description, generation, and interaction of eddies with the mean ocean circulation were topics that stirred physical oceanographers over the next decade while new instruments and techniques were being developed. As Henry Stommel put it, "Ideas have their own natural time to become feasible. MODE was an experiment that had an eight-year gestation period before it really could be done."

First, scientists deduced some tentative features of eddies by looking over existing hydrographic data. Evidence was found for eddies of differing shapes, sizes, and ages in the Atlantic, Pacific, and Arctic oceans. They appeared at all depths with widths of hundreds of kilometers and lifetimes of months or more. They moved clockwise or counterclockwise with colder or warmer cores of water and energy as much as 10 times greater than the sea circumscribing them.

Prevailing ocean circulation theory needed significant revision to include these energetic eddy motions. The goal of MODE scientists was to determine empirically the existence of deep-sea eddy motions and to place experimental data in a theoretical framework so that eddy dynamics could be included in open-ocean models (Figure 4).

Ready to Go

By 1970, the large-scale moored instrumentation program necessary for the MODE experiment was feasible. Experimentatlists and theoreticians worked together to define what measurements should be made and how they might be carried out. MODE became a test-bed for a diverse array of new measruing devices. Blending these different experimental approaches with theoretical studies would yield a three-dimensional view of an eddy.

The National Science Foundation's International Decade of Ocean Exploration (IDOE), together with the Office of Naval Research, were the major sources of scientific funding. Collective scientific authority and responsibility for the efforts of principal investigators representing

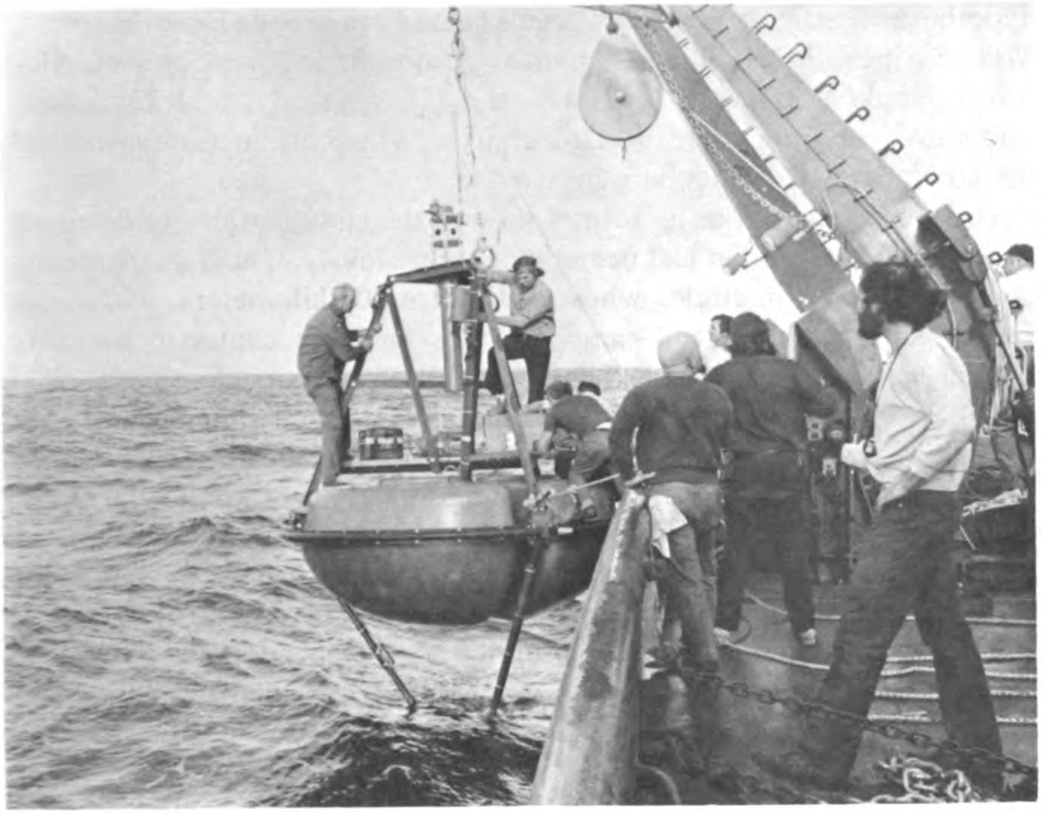

Figure 4 — Preparation of central marker buoy during MODE Project aboard CHAIN in 1973

15 oceanographic laboratories and institutions in the United States, England, Sweden, and West Germany resided in a 21-member MODE Scientific Council. Headed by Co-Chairmen Allan Robinson of Harvard and Henry Stommel along with Executive Officer Dennis Moore of Nova University, the Scientific Council and its various committees met regularly to exchange information, assess and interpret data, and make evaluations of program progress.

The Design

As a first step in the experimental/theoretical design of MODE, three arrays of current meters were set in the MODE area during 1971-1972 by the Woods Hole Oceanographic Institution's Moored Array Group. From these early current meter and hydrographic station results, the development and testing of numerical models and field statistics for the MODE area began. Theoretical studies, conducted during the 1972 summer-long MODE Theoretical Panel discussions, produced the design specifics for the experiment, evaluation of the accuracy of data types and sampling rates, techniques to map the temperature and salinity fields, numerical experimentation, and dynamical hypotheses. Participating theoreticians came from Woods Hole, Harvard, MIT, Scripps, the National Center for Atmospheric Research (NCAR), Johns Hopkins University (JHU), the Geophysical Fluid Dynamics Laboratory at Princeton, the University of Hamburg in Germany, and the University of Gothenburg in Sweden.

The preliminary-mooring results showed the eddy motions to be much larger horizontally than had been thought previously. Thus, instrumentation was laid out in circles whose radii were 100 kilometers, with more dense instrumentation and sampling rates near the center to measure smaller-scale eddy features. The site finally chosen centered on 28°N 60°40′W with a radius of 300 kilometers. The eastern half of the area is characterized by a rough, hilly region and the western half by the smooth Hatteras Abyssal Plain that slopes gently upwards to the Blake Plateau; thus, the effects of both rough and smooth topography on the eddy motions could be monitored. The site was accessible to east coast ports but sufficiently distant from the Gulf Stream to minimize the effects of its strong current. The field program was planned to last for four months, governed largely by the ship-time available and by instrument capability. That time period was believed sufficient to "catch" an eddy for ultimate description.

Field Program

In laying 17 moorings during two WHOI cruises, a total of about 75 tons of mooring hardware and instrumentation had gone over the
tons of mooring hardware and instrumentation had gone over the side. Bob Heinmiller, Chief Scientist on these legs and on the recovery leg four months later, remembers the logistical problems that had to be overcome in this, the largest effort to date by the WHOI Buoy Group: "In the same physical space where we used to prepare three or four moorings, we had to ready 17 moorings. We even had to rent space down the street at the Marine Biological Laboratory to prepare some moorings." At sea, two complete crews were designated for instrumentation launch so that work continued smoothly day and night. Heinmiller recalls, "As soon as one mooring was over the side, the crew was making preparations for the next one" (Figure 5).

Launching a mooring at a precise position is a delicate operation and an art in itself. While the ship steams upwind, the 5,500 meters of line with instrumentation attached is payed out from the stern. The entire mooring is towed until the designated position is reached. At that moment, the 2,550-pound anchor is released, and the mooring sinks into place. "The fact that we had laid all the moorings within a mile or so of the specified position was something we were proud of," Heinmiller adds. "In the past we considered ourselves fortunate to be within four

Figure 5 — Current meter ready to go overboard as part of mooring.

or five miles of the position." Although problems were encountered on two mooring launches, the crews were sufficiently prepared to deal with the situation smoothly. Heinmiller reflected on his 12-year involvement in buoy operations: "Though we still have problems to solve, the Buoy Group has graduated in the past five years from problems of survivability to problems of performance."

However, the moored instrumentation effort was but one of 12 programs simultaneously carried out over the four-month field program. To better understand the scope of the MODE-1 field program, imagine a large column of water which has a diameter of 600 kilometers and a depth of five kilometers. In that column of water 25 moorings have been implanted with over 90 current meters and 60 temperature-pressure recorders strategically placed at several depths throughout the column. Moving horizontally through the water at approximately 1,500 meters from the surface are 20 SOFAR floats whose paths are recorded acoustically at land-based hydrophone stations. Also moving horizontally through the water at 600-700 meters from the surface is a ship-board-controlled towfish which is following a selected density surface. At intermittent intervals averaging seven times daily, instruments are moving vertically through the water column: a total of 800 deep hydrographic stations are made at 77 fixed grid points; 50 electromagnetic and 11 acoustic profiles (measuring the changing velocity and direction of the water column) are taken; and more than 1,300 air-dropped probes (measuring the horizontal transport of the water) are released during AirMODE's 46 flights. Implanted just above or on the sea floor are some 20 different instruments measuring the magnetic and electric field fluctuations and the variations in temperature and pressure. On the eastern side of the column is the MiniMODE site. Four moorings with five current meters on each are implanted in this more rugged topographic zone. During April and May, up to 22 floats are simultaneously tracked inside a dense pattern of hydrographic stations.

The Results

When the last ship pulled out of MODE waters in mid-July, 1973, scientists and technicians returned home to face the massive pile of data collected during the experiment. The raw data required long hours of evaluation, analysis, interpretation, and integration into the MODE data set. During the summer of 1974, MODE scientists participated in a two-month Summer Institute at the University of Rhode Island (URI) for purposes of generating a synoptic atlas of the data, an intercomparison of measurements from different instruments, and a theoretical dynamical discussion of the results.

Their data showed that an eddy had passed right through the center of the MODE site, with evidence of other eddies surrounding it. The central eddy had a warm core, clockwise motion, and drifted slowly westward at about two kilometers per day. Then a new eddy was seen sweeping in from the east, bringing with it a cold-water core and counterclockwise motion. From horizontal and vertical profiles collected, the current fields in these eddies were seen to parallel lines of constant temperature, a feature which was also reflected in Swallow's free-drifting floats. However, the trajectories of the 3,000-meter-depth Swallow floats and the 1,500-meter depth SOFAR floats suggest topographic influence. The vertical profiling measurements also indicated that bottom topography affects the vertical distribution of currents: a typical profile taken over the central smooth area was significantly different from those over the rough eastern area. Sharp temperature differences recorded between the warm surface layers of the ocean and the cold deeper layers may contain a source of energy for the eddies. Mechanisms which operate to generate eddies that, taken in some combination, may hold the key to the eddy dynamics also include boundary conditions and atmospheric forcing.

What Next?

Now that scientists have some idea of the character of an eddy, they are looking into the how's and why's of its existence. Where are eddies formed, where do they go, and why? How do they interact among themselves and with the large-scale circulation? These questions will be addressed by U.S. and U.S.S.R. scientists in the POLYMODE experiment.

Discussions with a 10-member U.S.S.R. delegation at the 1974 Summer Institute offered valuable input for the analysis of MODE-1 data in light of results from the 1970 U.S.S.R. Polygon experiment that also studied eddy phenomena. The Joint Agreement for Scientific and Technical Cooperation, signed by President Nixon and Secretary Brezhnev in 1972, was instrumental in securing American and Russian participation in a cooperative program for oceanographic research.

The design for POLYMODE is in an evolutionary stage, but it will be a longer, more extensively instrumented program covering a larger swath of the Atlantic Ocean. The 1976-78 field program will include a more-intensive survey period for the 12 months of 1977. Current meters, SOFAR floats, CTD's, and other instrumentation will be deployed.

The enigmatic eddies will continue to challenge the imagination of scientists for years to come. Since ocean currents carry large amounts of heat, chemical substances, and planktonic organisms about the earth, variations in the currents caused by eddies, as well as other oceanic phenomena, can have major effects on weather patterns, fishing, pollution dispersion, and coastal recreation. Predicting the eddies' movements within, and their effects upon, the world's ocean patterns are long-range goals of many physical oceanographers. MODE and POLYMODE are only a beginning.

Mosquito Control with Monomolecular Surface Films

A novel form of mosquito control has been proved highly effective in both laboratory and field experiments by W. D. Garrett of the Ocean Sciences Division of the Naval Research Laboratory. The approach uses nonionic surface-active liquids which spread spontaneously over the surface of water and form films which are nearly monomolecular (one molecule in thickness). The quanity of material required for the control of mosquito larvae and pupae in a body of water is several orders of magnitude less than the amounts of petroleum-based oils presently used. In contract to the petroleum oils, it is likely that these thin nontoxic films will be suitable for use on reservoirs and irrigation water bodies. Furthermore, they do not interfere with gas exchange across the air-water interface, and thus will not cause oxygen starvation in the underlying water.

The efficacy of this approach to mosquito control was recently demonstrated in laboratory and field experiments at the Jacksonville Naval Air Station. The enthusiastic cooperation and entomological expertise of the personnel of the Navy Disease Vector Ecology and Control Center were instrumental in the accomplishment of this research. Experiments were performed on fourth-stage larvae reared for these studies by the USDA, Insects Affecting Man Research Laboratory, Gainsville, Florida. The larvae were *Anopheles quadrimaculatus*, a primary malaria carrier. In the laboratory, several of the organic films caused 100% mortality in fourth-stage larvae in thirty minutes. One of the films was nearly 100% fatal to the pupae in 75 minutes. Field studies were performed in square ponds with an area of 4 square meters each in a natural paludal setting, with two of the ponds used as controls. It was possible to maintain effective surface films for at least 24 hours from a small initial application. Mortalities to larvae were essentially the same in the field as had been determined in the laboratory, although it was difficult to determine the exact time at which total mortality was effected. Two of the chemical films were 100% effective against larvae in the field studies, each experiment involving about 5000 organisms.

Although the mechanisms by which the films cause mortality in aquatic forms of the mosquito are not exactly known, the effects are physical rather than toxic in nature. The greatly reduced surface tension caused by the monomolecular layers prevents mosquito larvae from attaching themselves in a proper breathing orientation at the water surface. Pupae appear to have difficulty penetrating the organic film, and it is also possible that the constitutents of the film increase the wettability of the internal surface structures of pupal trumpets and larval siphons (surface breathing tubes) allowing water to interfere with respiration by the organisms. The greatly reduced surface tension may sink mosquito eggs of species which require flotation for proper hatching, interfere with adult females attempting to deposit floating eggs, and drown adults emerging from their pupal shells.

Automatic surface-film dispensers were used to maintain organic film coverage at high surface pressures (low surface tension) for extended periods of time. However, a single application at a concentration of 40 mg per square meter was adequate to maintain effective coverage for one day. This was possible because of the self-repairing properties inherent in fluid surface films of the type used in this research. The excess liquid mosquito-control agent remains as an unspread, oil-like lens in equilibrium with its own film. Thus, the excess liquid acts as a reservoir to rapidly replace any loss in the organic film, thereby maintaining the pressure in the compressed monomolecular film.

Since the use of DDT has been terminated in this country, and the manufacture of other pest-control chemicals such as aldrin, and dieldrin has been greatly curtailed, alternatives to chemical pesticides are urgently needed. The surface film approach to mosquito control is based on the modification of physical processes necessary to the lift-style of the mosquito in water. The mosquito is killed prior to its emergence into the atmosphere as an adult. This technique promises to be a harmless and uncomplicated means of controlling one of the world's major disease vectors—the mosquito.

Unique Energy-Storage System

Scientists at the Naval Research Laboratory (NRL), have come up with a unique energy-storage system.

The investigators say their system stores mechanical energy in high-speed flywheels and then extracts it in electrical form by applying a magnetic field parallel to the axis of the wheels; a radial current is generated which charges the storage inductor. The electrical energy is then rapidly transferred to a load through the opening of the inductive circuit.

Drs. Anthony Robson and William Lupton, principal investigators for the NRL program report that a unique feature of their system (a Faraday disc, or homopolar generator) is that the excitation coil of the generator is also the energy-storage coil, making it an "inertial-inductive storage system" which is extremely compact.

Until now, Robson and Lupton report, capacitor banks have provided the principal source of pulses power for a variety of systems such as lasers, electron-beam generators and high-energy-density plasma generators for high-intensity radiation sources or pulsed thermonuclear reactors.

But as the energy requirements increase above the level of a few megajoules (MJ), capacitor banks become large, cumbersome and costly. The potential of the NRL system meets the growing need for compact pulsed-power sources. Nominally rated at 10 MJ, the NRL system occupies about 200 cubic feet. A capacitor bank storing the same energy would occupy at least 20,000 cubic feet.

This system is currently operating at a level at about one MJ. Eventually, it should transfer about 5 MJ's a load in 500 micro-seconds, which represent a pulsed-power output of 10 gigawatts (GW). Power supplies of several hundred megajoules could be constructed with series-parallel combinations of the basic 10-MJ module.

NAVAL RESEARCH REVIEWS publishes highlights of research conducted by Navy laboratories and contractors and describes important naval experimental facilities. Manuscripts submitted for publication, correspondence concerning prospective articles, and changes of address, should be directed to Code 730, Office of Naval Research, Arlington, Va. 22217. Requests for subscriptions should be directed to the Superintendent of Documents, U.S. Government Printing Office, Washington, D.C., 20402. Subscription price: $6.50 per year in the U.S. and Canada; $8.15 per year, foreign; $0.55 per individual copy. The issuance of this periodical approved in accordance with Dept. of the Navy publications and printing regulations.

NAVSO P-510

IN THIS ISSUE

VOL. XXIX, NO. 4

Cover Caption

Launching a mooring at a precise position is a delicate operation and an art in itself. *See article on page 22.*

DEPARTMENT OF THE NAVY
OFFICE OF NAVAL RESEARCH
ARLINGTON, VA. 22217

OFFICIAL BUSINESS

PENALTY FOR PRIVATE USE, $300.

POSTABE AND FEES PAID
DEPARTMENT OF THE NAVY
DoD-316

U.S.MAIL