


D 210.11:
43/2

Naval Research Reviews

Office of Naval Research
Two / 1991
Volume XLIII





A meteorologist releases a balloon for wind profile measurement while another tracks it with a surveying instrument in the 1970's. Future tracking will be done by the Global Positioning System (GPS).

The Arctic Ice Dynamics Joint Experiment (AIDJEX)

AIDJEX was a cooperative U.S. and Canadian research venture in the Arctic Ocean to better understand the large-scale response of sea ice to its environment. In 1969 the Office of Naval Research (ONR) contracted Dr. Kenneth Hunkins of Lamont Doherty Geophysical Laboratory of Columbia University and Dr. Norbert Untersteiner of the University of Washington to plan a multi-station gathering complex which consisted of four manned ice island stations surrounded by a circle of 15 automatic, battery powered stations (data buoys). A series of pilot studies began in 1970 to resolve scientific and technical questions before the year long data gathering program began in 1975.

AIDJEX was a coordinated program of data acquisition and mathematical modeling. Its specific purpose was to find a quantitative relationship between large-scale stress and strain fields in sea ice with methods to determine the external stresses exerted on the ice by wind and water currents. These findings led to estimates in forecasting ice convergence, divergence, or shear, which information is necessary for off-shore drilling and surface shipping in ice covered seas.

The dynamic ice model developed by AIDJEX was an important step toward understanding the long-term climatic interaction between the atmosphere, cryosphere, and hydrosphere as well as providing a basis for an analysis of the role of ice-covered seas in world climate.

Naval Research Reviews

Office of Naval Research
Two / 1991
Volume XLIII

Articles

2

Identifying Ice Floes in Satellite Images

*Jeffrey D. Banfield
Adrian E. Raftery*

13

Neural Networks: A Computer Science Perspective

Dr. Alan L. Meyrowitz

19

Intelligent Tutoring Technology: The Emergence of Individualized Instruction

Bruce W. Hamill

DEPOSITORY

AUG 27 1991

UNIVERSITY OF ILLINOIS
AT URBANA-CHAMPAIGN



CHIEF OF NAVAL RESEARCH
RADM William C. Miller, USN

DIRECTOR
Dr. Fred Saalfeld

CHIEF WRITER/EDITOR
William J. Lescure

SCIENTIFIC EDITORS
*Dr. Robert J. Nowak
Dr. J. Dale Bultman
Sarah Scott*

MANAGING EDITOR
Norma Gerbozy

ART DIRECTION
*Typography and Design
Desktop Printer, Arnold, MD*

Departments

12

Profiles in Science

25

Research Notes

About the Cover

A satellite image in the infrared taken by an operational line scanner of an ice floe in Baffin Bay off the coast of Greenland. Photo is courtesy of the Remote Sensing Division of the Naval Oceanographic Atmospheric Research Laboratory. (See article on page 2)

Naval Research Reviews publishes articles about research conducted by the laboratories and contractors of the Office of Naval Research and describes important naval experimental activities. Manuscripts submitted for publication, correspondence concerning prospective articles, and changes of address, should be directed to Code OPARI, Office of Naval Research, Arlington, VA 22217-5000. Requests for subscriptions should be directed to the Superintendent of Documents, U.S. Government Printing Office, Washington, DC 20403. *Naval Research Reviews* is published from appropriated funds by authority of the Office of Naval Research in accordance with Navy Publications and Printing Regulations. NAVSO P-35.

Original from
UNIVERSITY OF ILLINOIS AT URBANA-CHAMPAIGN
Two/1991 1

Identifying Ice Floes in Satellite Images

Jeffrey D. Banfield
Montana State University

Adrian E. Raftery
University of Washington

Abstract

Identification of ice floes and their outlines in satellite images is important for understanding physical processes in the polar regions, for transportation in ice-covered seas and for the design of offshore structures intended to survive in the presence of ice. At present this is done manually, a long and tedious process which precludes full use of the great volume of relevant images now available.

We describe an automatic and accurate method for identifying ice floes and their outlines. Floe outlines are modeled as closed principal curves, a flexible class of smooth non-parametric curves. Initial estimates of floe outlines come from the erosion-propagation (EP) algorithm, which combines the idea of erosion from mathematical morphology with that of local propagation of information about floe edges.

The edge pixels from the EP algorithm are grouped into floe outlines using a new clustering algorithm. This extends existing clustering methods by allowing groups to be centered about arbitrary curves rather than points or lines. This may open the way to efficient feature extraction using cluster analysis in images more generally. The method is implemented in an object-oriented programming environment for which it is well suited, and is quite computationally efficient.

1. Introduction


Knowledge of the shapes, sizes and spatial distribution of ice floes is important for understanding the physical processes operating on the ice pack in the polar regions. It is also important for practical problems associated with transportation in ice-covered seas and for the design of offshore structures intended to survive in the presence of ice.

Such information can be found in satellite images of the polar regions such as Figure 1, which exist in large and rapidly increasing numbers. Practical use of such images requires identification of the outlines of ice floes above a certain size. To date this has been done manually (Rothrock and Thorndike, 1984), a slow and tedious process that often takes a day or more to record the data from a single image and effectively precludes full use of the data. Automating the process is inherently difficult. Problems include the presence of many smaller floes and of melt ponds on the surface of floes which ensure that floes often do not appear as homogeneous blocks of ice in the image.

In this article we describe an automatic method for identifying the outlines of ice floes. The outcome of this is shown in Figure 2, and is almost the same as the result of very careful manual digitization. We model ice floe outlines as closed

Figure 1

A polar LANDSAT image showing ice floes. This is a 200 x 200 pixel image, where each pixel is 80m square; it thus represents a 15 x 15 km area.



principal curves (Hastie and Stuetzle, 1989—hereafter HS), a flexible family of one-dimensional non-parametric curves in a higher-dimensional space. Our method consists of identifying a set of edge pixels and grouping them into clusters about a principal curve. Each cluster corresponds to a floe and the corresponding principal curve is the estimated floe outline.

The method involves several new statistical techniques which have been developed for this application:

- (1) A way of estimating closed principal curves that reduces both bias and variance and is robust to outliers. Here, outliers take the form of melt ponds on the surface of ice floes. Principal curves are non-parametric curvilinear analogues of principal components that minimize the sum of squared projected distances from a set of data to the curve. Here, they are used to model the outlines of the ice floes.
- (2) The erosion-propagation (EP) algorithm provides initial estimates of floe outlines. This combines the existing idea of erosion from mathematical morphology with that of local propagation of information about floe boundaries.
- (3) A method for clustering about principal curves. Existing clustering algorithms separate data into groups, each of which is clustered about some central point. Here we generalize this to allow each group to be clustered about a different curve. This opens the possibility that cluster

analysis may be useful more generally for fast curvilinear feature extraction in images.

The method is implemented in an object-oriented programming environment for which it is well suited, and seems computationally efficient.

2. Principal Curves

In this section, we first review the definition and basic properties of principal curves (Section 2.1). We then describe a new robust unbiased algorithm for estimating closed principal curves. (Section 2.2).

2.1 Non-parametric Smoothing and Principal curves

Consider a data set consisting of n measurements made on two variables, x and y , as shown in Figure 3, and the problem of trying to summarize the relationship between x and y . When y is *dependent* on x , so that, for example, $y=f(x)+\epsilon$, where ϵ is a random variable with mean zero, the joint relationship between x and y can be described using regression. It is commonly assumed that $f(x)$ is linear, but this is often inappropriate. A flexible approach is to proceed non-parametrically without assuming any prespecified functional form for $f(x)$. This leads to the idea of non-parametric regression where

Figure 2

The ice floe outlines, larger than a fixed minimum size, found by our procedure for the data in Figure 1.





Figure 3.

An illustration of the difference between the first principal component (the solid line) and a principal curve (the dashed line). The first principal component is the straight line that minimizes the sum of squares projected distances from the data to the line. The principal curve is a non-parametric analog of the first principal component. It is a "smooth" curve that minimizes the sum of squared projected distances from the data to the curve. The principal curve can adapt to the shape of the data.



$f(x)$ is assumed to be a smooth non-parametric curve which may be estimated, for example, using a spline or kernel smoother. Non-parametric regression procedures, often called smoothers, smooth out the variability in the data and provide a non-parametric estimate of $f(x)$.

In the symmetric situation, where one variable does not depend on the other, the usual linear summary of the joint behavior of x and y is provided by principal components. The first principal component is the line that minimizes the sum of the squared *projected* distances from the data points to the line. A principal curve is a non-parametric analogue of the first principal component. It is a smooth one-dimensional curve that passes through the middle of a data set in such a way as to minimize the sum of the squared projected distances from the data to the curve. Figure 3 shows a set of data together with the first principal component and a principal curve.

Suppose that we wish to predict the value of f at a point x^* . The assumption behind smoothers is that values of y associated with values of x near x^* should be "near" $f(x^*)$. Thus, by averaging the values of the y 's associated with x 's in a neighborhood of x^* we can generate a non-parametric point estimate of $f(x^*)$ and, by repeating this procedure over the entire range of x 's, we can generate a non-parametric estimate of f .

2.2 Estimating Principal Curves


In non-parametric regression the dependency between x and y makes the smoothing process relatively simple: the neighborhoods are defined on the x 's and the smoothing is performed on the y 's. In the estimation of principal curves the two variables are treated symmetrically, so that neighborhoods cannot be defined using just one of the variables, and it is necessary to smooth both variables using the arc-length of the principal curve to define neighborhoods.

In order to estimate a principal curve we need to define a smooth one-dimensional curve, $f(\lambda)$, as a vector-valued function consisting of an x -component and a y -component, both parameterized by the single parameter λ , which is often the arc-length of the curve. The neighborhoods are defined by λ , that is, points that have projections near each other on the curve are considered neighbors no matter how far apart they may be in the plane. The smoothing is componentwise: the x -components of points in the same neighborhood are averaged, and so are the y -components, yielding a smoothed x -component and a smoothed y -component. We are actually performing two non-parametric regressions, where the x -component and the y -component are individually regressed on λ .

Principal curve estimation is an iterative procedure, each iteration consisting of two steps. In the first step the data are

Figure 4.

Estimated principal curves for simulated data. The data were obtained by generating points uniformly on the circumference of a circle, and perturbing them randomly along the normal to the circle according to a Gaussian distribution. The principal curves were estimated using the algorithm (2.2) of HS with spans of 0.2 (outer dashed line), 0.3 (inner dashed line) and 0.5 (solid line).



projected onto $\hat{f}_i$, the i^{th} estimate of the principal curve, and then ordered according to their projections. The ordering defines neighborhoods that are used in the coordinatewise smoothing of the second step to produce the next estimate, $\hat{f}_{i+1}$, of the principal curve. The data are then projected onto $\hat{f}_{i+1}$, and the process is repeated until the change between successive estimates of the curve falls below some threshold level. By projecting the data onto the new estimate of the principal curve at the start of each iteration the neighborhoods are allowed to change. Two points that were neighbors at the i^{th} iteration may not be neighbors at the $(i+1)^{\text{th}}$ iteration because they project to points that are far apart as measured by λ ; this allows the principal curve estimate to adapt to the shape of the data. By using projections, rather than vertical distances as in regression, the variables are treated symmetrically and the estimation procedure is invariant to rotation. This is a desirable property in a curve that will be used to model ice floes since we want the same floe outline regardless of the orientation of the image.


Smoothers generally produce curves that are biased towards the center of curvature. For closed curves such as the

one shown in Figure 4, the center of curvature is interior to the curve (except for small local regions in non-convex curves), often resulting in large biases. It is clear from Figure 4 that the estimated principal curve suffers from this problem, and that the bias increases with the smoothness of the curve. In most statistical problems there is a tradeoff between smoothness and bias on one hand, and variance on the other. Rather than smoothing the x - and y - components of the data at each iteration we modified the original HS algorithm by smoothing the residuals, or projection vectors from the data to the curve, and then adjusting the current estimate of the curve by the smoothed residuals. Figures 5 and 6 compare this algorithm with that of HS for various amounts of smoothing; it is clear that the bias problem has been solved.

Occasionally some of the data points do not belong to the main body of the data that defines the principal curve; these are called outliers and they can cause problems in the estimation procedure. In the satellite images that we analyze, outliers often arise in the form of shallow but sometimes large melt ponds on the surface of the ice floe. An example of this is Figure 7, which shows the edge pixels of one floe in Figure 1, as identified by the EP algorithm. The edge pixels near the left side interior to the floe are from melt ponds. Since they do not

Figure 5.

Estimated principal curves using the unbiased algorithm proposed in this article. This shows the principal curves resulting from our unbiased estimation algorithm on the same data and at the same spans (.2, .3 and .5) used for the HS algorithm in Figure 3. The three curves almost totally overlay each other.



belong to the edge of the floe, we need to ensure that they not affect the estimate of the principal curve.

To eliminate the effect of outliers we used a weighted average in the smoothing step, where the weight of a point depends on its distance from the current estimate of the principal curve. We calculated a robust analogue of the standard deviation of the lengths of the projections (equal to 1.25 times the mean absolute deviation). We then set the weight for a point to zero if it is more than three robust standard deviations from the current estimate of the principal curve, and to unity otherwise. Figure 7 shows the result of this robust procedure as well as that of a non-robust procedure which uses the mean of all the data in each neighborhood. The robust procedure has clearly achieved its goal.

3. The Erosion-Propagation (EP) Algorithm

To select the potential edge pixels and provide an initial grouping of them into floe-outlines, we used the EP algorithm. This operates on binary images. However, images of ice floes such as Figure 1 are usually greyscale. The marginal distribu-

tion of pixel intensities is highly bimodal, and so we work with the simpler binary image obtained by thresholding the original image; see Figure 8. The final result is relatively insensitive to the precise choice of threshold.

The erosion part of the EP algorithm, which identifies the potential edge elements, is a standard application of ideas in mathematical morphology (Serra, 1982). The propagation part of the EP algorithm keeps track of the floe to which an edge pixel belongs by locally propagating the information about edge elements into the interior of the floe as it is eroded. This is facilitated by the object-oriented programming environment.

The algorithm is iterative and operates on a binary image consisting of figures (ice floes) on a contrasting background (water). At the first iteration, if a pixel is ice and a specified sub-set of its neighbors is water, the pixel "melts" and becomes water; in this way the figure to which it belongs is eroded. In our implementation, a pixel "melts" if any of the eight neighboring pixels are water. At the second iteration, the same operation is performed on the image resulting from the first iteration, and so on. This can be formally described in terms of structuring elements using the terminology of mathematical

Figure 6.

The principal curve from the HS algorithm (dashed line) at a span small enough to eliminate most of the bias, compared with the estimate proposed here (solid line) with a span of 0.2. The smaller the span, the less the bias and the rougher the estimated curve. Notice how rough the HS curve is, compared with our estimate.




Figure 7.

The small circles are the edge pixels for one of the floes in Figure 1, as identified by the EP algorithm. The points interior to the floe are from melt ponds. The lines show a principal curve estimated using the robust procedure described in the text (thick line), compared with the estimate from the non-robust procedure (thin line). The robust estimate is unaffected by the melt ponds, while the non-robust estimate is pulled towards them.





Figure 8.

Binary version of Figure 1 for two threshold levels. The results are similar. However, (A) has a lower threshold level than (B) and therefore has more clutter in the water but less noise interior to the floes. The results of applying the EP algorithm to (A) is shown in (C) after three iterations and (D) after 12 iterations.



morphology (Banfield, 1988). The potential edge pixels are those that are eroded at the first iteration of the EP algorithm.


Some results are shown in Figure 8. We can control the minimum size of the floes by waiting until a specified number of iterations, $i_{\min}$, have passed before recording a floe. The smallest floe which can be recorded is then a square of side $(2i_{\min} + 1)$ pixels. Smaller floes “melt” and are not recorded.

The idea of the propagation part of the EP algorithm is that the locations of the edge pixels are propagated towards the interior of the figure as it is eroded. At the end of the process, a single interior point of the figure will “know” the locations of all the edge pixels to which it corresponds. The location information is passed to only a few pixels which are taken as far from the eroded pixel as possible subject to them not belonging to a different floe. This ensures that the amount of location information to be processed does not become unmanageable. It also prevents loss of information due to irregularity of the floe, melt ponds, or pixel misclassification at the thresholding stage. All of the eroded pixels are processed in exactly the same manner and it is this uniform processing that allows the algorithm to be implemented on parallel processing machines.

In Figure 9 we show the results of the EP algorithm applied to the data in Figure 1 with a minimum floe size of 15×15 pixels (i.e. 1.2 km. square). The results are reasonably

Figure 9.

Result of the EP algorithm applied to the data in Figure 1. Floes are not recorded unless they have survived at least 7 iterations. This corresponds to a minimum floe size of 15×15 pixels, or 1.2 km square. The open circles are the edge elements identified by the EP algorithm. The numbers (or solid dots) interior to each floe are the centers found by the EP algorithm. Note that centers 1 and 4 are on the same floe, which was subdivided because of the melt ponds. Other floes were also subdivided.



good: of the 35 floes identified by the Ep algorithm, 23 are “right” in the sense of being close to the floes identified by careful manual digitization.

However, the EP algorithm tends to subdivide floes. This can occur when the floes are non-convex or when they have noise in the interior. Figure 10 shows an example of the non-convex case. As the floe shown in Figure 10 was eroded, the narrow middle section was pinched in and the floe was divided into two partial floes. Figure 11 is an example of a floe with melt ponds that cause the EP algorithm to produce three partial floes.

4. Clustering About Closed Principal Curves


The EP algorithm tends to subdivide floes. We have therefore developed a method for determining which of the floes identified by the EP algorithm should be merged, based on an algorithm for clustering about closed principal curves.

Original from

UNIVERSITY OF ILLINOIS
URBANA-CHAMPAIGN

Figure 10.

This figure shows how, when a non-convex floe is eroded, the narrow region can be "pinched off" resulting in two partial floes (indicated by a "+" and a "o")



Since we want to find out whether to merge tentatively identified floes, this is hierarchical and agglomerative.

The objective of cluster analysis is to group a set of observations into "interesting" sub-sets. In practice, this has usually meant grouping observations which are close to one another. Ward (1963) proposed a hierarchical agglomerative algorithm for dividing data into g groups such that the sum of the within-group sums of squares is minimized. The algorithm starts by assigning each observation to a separate group. At each agglomeration two groups are merged, chosen so as to minimize the increase in the sum of within-group sums of squares. This clustering criterion is optimal if the data are generated by a finite mixture of spherical normal distributions. This corresponds to clusters which tend to be of the same size and spherical.

Murtagh and Raftery (1984) proposed decomposing the within-group sum of squares into parts and using a weighted sum of the parts as the clustering criterion, with weights chosen so as to emphasize aspects of interest in the application. For two-dimensional data, they suggested decomposing the within-group sum of squares into parts parallel and perpendicular to the first principal component of the group and downweighting the parallel part. This criterion was generalized by Banfield and Raftery (1989) who also showed that it

is optimal when the data are generated by a mixture of normal distributions with covariance matrices whose eigenvalues are constant across clusters. This corresponds to clusters which tend to be elliptical with the same size and shape but different orientations.

We now apply the idea of decomposing and reweighting the within-group sum of squares to the present problem. The edge pixels for an ice floe that has not been subdivided should be

(a) tightly clustered about the floe outline, as estimated by the principal curve, and

(b) regularly spaced along the outline, so that the variance of the distances between neighboring edge pixels should be small; see Figure 12.

The characteristic that allows the successful merger of the potential floes is that the edge elements of a complete floe will have evenly distributed projections onto the principal curve while large gaps in the projections indicate a partial floe. A clustering criterion of the form

$$V^* = \alpha V_{\text{about}} + V_{\text{along}} \quad (4.1)$$

provides an adjustable measure of how the edge elements project onto the principal curve. In (4.1), V_{about} is the sum of squares of the lengths of the projections of the data onto the

Figure 11.

Noise in the interior of a floe can erode outwards and cause the floe to be subdivided. In this case, the melt ponds in the center caused the floe to be subdivided into three partial floes (indicated by an "x", "+" and a "o")






Figure 12.

A complete floe (A) and a partial floe (B), showing the edge pixels identified by the EP algorithm (open circles) and their projections onto the estimated principal curve (solid circles). The distance between adjacent projections in (A) has a smaller variance than the distance between adjacent projections for the partial floe shown in (B).




A



B

Figure 13.

The circles are the floe edge elements found by the EP algorithm. The lines are the principal curves of the floes which were estimated after using the clustering method to merge partial floes.



principal curve and V_{along} is the sum of squares of the distances along the principal curve between adjacent projections.

To determine whether a set of floes should be merged we calculated V^* for each of the individual floes. We then calculated V^* for the floe that would result from the merger of the edge elements of the individual floes. If the floe resulting from the merger has a smaller value of V^* than any of the individual floes, the merger is needed. Otherwise, the individual floes should not be merged.

To determine α , we noted that, by arguments similar to those of Banfield and Raftery (1989), V^* will be an optimal criterion, conditional on the estimated principal curves, if the edge pixels are normally distributed about the floe outlines and if $E/V_{\text{along}} = \alpha E/V_{\text{about}}$ where E denotes statistical expectation. We therefore estimated α as the average of $V_{\text{along}}/V_{\text{about}}$ for the floes that we knew were not subdivided, namely those which had no shared edge elements. Using a neighborhood defined by the nearest 30% of the data to estimate the principal curves yielded $\alpha = 0.39$. Reasonable choices for the amount of smoothing can be determined by the time of year (early summer floes are rough and jagged while late summer floes have lost their rough edges) and location (open ocean, marginal ice zone or within the ice pack). The procedure is not unduly influenced by the choice of the smoothing parameter

Original from

UNIVERSITY OF ILLINOIS
URBANA-CHAMPAIGN

and, since $\hat{\alpha}$ is estimated from the complete floes, the procedure can adapt to whatever smoothing parameter is used.

The method gives results that are both correct and clear-cut. Figure 13 shows the final results of the procedure, together with the identified edge pixels.

5. Discussion

We have described an automatic method for finding the outlines of ice floes in satellite images. It is accurate and computationally efficient. It involves three new statistical techniques: a robust method for estimating closed principal curves, the EP algorithm, and a method for clustering about principal curves. A fuller and more detailed technical description of the methods is contained in Banfield and Raftery (1991).

The approach would seem to be applicable more generally to the detection of non-linear features in images. It extends cluster analysis to the case where similar pixels tend to be grouped about arbitrary curved features, open or closed, using the idea of decomposing and reweighting the within-group sum of squares proposed by Murtagh and Raftery (1984). This suggests that cluster analysis may be useful for feature extraction in images more generally.

The procedure is implemented in an object-oriented programming environment. One of the advantages of this environment is that each floe resulting from the procedure can be represented as an instance of a "floe object" and can carry with it information about the floe to be used in subsequent stages of the analysis. It is relatively fast: a 512x512 8-bit image can be analyzed in about three minutes on a Sun SPARCstation 1. The processing time is linear in the number of pixels, but does depend upon the complexity of the image. The EP algorithm has the potential of being implemented on parallel processing machines.

To date, the development of automated techniques for the analysis of polar satellite images has been limited to ice floe tracking (Ninnis, Emery and Collins 1986; Fily and Rothrock 1986, 1987; Vesecky, Samadani, Smith, Daida and Bracewell 1988). The primary tool in these automated tracking methods is cross-correlation, which provides the ability to match regions in two different images, but does not give any information about the morphology of the individual ice floes or the spatial structure of the ice pack. Vesecky *et al.* (1988) use segments of ice floe boundaries to track ice floe movements, but this does not provide the type of information needed to study ice floe morphology and spatial structure. The need for more information on both morphology and spatial structure was clearly shown by the 1984 Marginal Ice Zone Experiment (Burns *et al.* 1987; Campbell *et al.* 1987). Floe outlines estimated using the methods described here have been used to track individual floes through a series of remotely sensed images (Banfield 1991).

The problem considered here does not fall neatly into one of the problem areas in image understanding that have been

intensely studied in recent years, namely image restoration, classification, segmentation and feature extraction. It does, however, combine elements from all of them. Image restoration attempts to reconstruct a degraded image. Image classification tries to assign each pixel to one of several predetermined categories; it may be regarded as a special case of restoration. Image segmentation (Rosenfeld and Kak, 1982) seeks to identify areas of contiguous pixels that are, for example, devoted to the same crop. The aim of feature extraction is to find linear or curvilinear features in images.

Our problem shares goals with feature extraction, segmentation and classification. While restoration is not an explicit goal, we would expect the methods developed here to work well in the presence of degradation. Restoration and classification methods do not, by themselves, address the present problem. Current feature extractors would seem to have difficulty locating features as arbitrary as ice floes. For example, the generalized Hough transform (Ballard 1981), an obvious candidate for locating closed curves, requires an initial pattern description which it then tries to find in the image. It would be difficult to provide an initial pattern description that is general enough to accommodate the wide range of commonly found ice floe shapes. Our approach may also be applicable to segmentation problems, especially those concerned with identifying not only regions but also the shapes of their outlines.

The Bayesian and stochastic relaxation approach of Geman and Geman (1984) may well be applicable to the present problem, although it has to date been used mainly for restoration. It would require extensive modeling assumptions for the ice floe problem, and experience suggests that it would be computationally expensive. The computational burden might be reduced by using as an approximation the ICM algorithm of Besag (1986), although this does not yet appear to have been applied to problems such as the present one. Our procedure on the other hand requires only the assumption that the ice floe boundaries be closed curves, and it is relatively fast.

Biographies

Jeffrey D. Banfield is Assistant Professor, Department of Mathematical Sciences, Montana State University, Bozeman, Montana. His publications and interests are in statistical computing and developing new statistical methodology for use in pattern recognition and remote sensing problems.

Adrian E. Raftery is Professor of Statistics and of Sociology, University of Washington, Seattle, Washington. He has written over thirty articles in scientific journals on spatial statistics, time series analysis, Bayesian methods, and the development of new statistical methods for scientific applications in sociology, biology and the environmental sciences. He is associate editor of the *Journal of the American Statistical Association* and the *Journal of the Royal Statistical Society*.

References

1. Banfield, J. (1988), "Constrained Cluster Analysis and Image Understanding," Unpublished Ph.D. Dissertation, Department of Statistics, University of Washington.
2. Banfield, J. (1991), "Automated Tracking of Ice Floes: A Stochastic Approach," *IEEE Transaction in GeoScience and Remote Sensing* (in press).
3. Banfield, J.D. and Raftery, A.E. (1991), "Ice Floe Identification in Satellite Images Using Mathematical Morphology and Clustering About Principal Curves," *Journal of American Statistical Association* (in press).
4. Banfield, J.D., and Raftery, A.E. (1989), "Model-based Gaussian and non-Gaussian clustering," Technical Report no. 186, Department of Statistics, University of Washington.
5. Ballard, D.H. (1981) "Generalizing the Hough Transform," *Pattern Recognition*, **13**, pp. 111-122.
6. Besag, J. (1986). "On the Statistical Analysis of Dirty Pictures (with Discussion)," *Journal of the Royal Statistical Society, series B*, **48**, 259-302.
7. Burns, B.A., Cavaleri, D.J., Keller, M.R., Campbell, W.J., Grenfell, T.C., Maykut, G.A. and Gloersen, P. (1987), "Multisensor Comparison of Ice Concentration Estimates in the Marginal Ice Zone", *Journal of Geophysical Research*, **92**, 6843-6856.
8. Campbell, W.J., Gloersen, P., Josberger, E.G., Johannessen, P.S., Guest, P.S., Mognard, N., Shuchman, R., Burns, B.A., Lannelongue, N. and Davidson, K.L. (1987), "Variations of Mesoscale and Large Scale Sea Ice Morphology in the 1984 Marginal Ice Zone Experiment as Observed by Microwave Remote Sensing," *Journal of Geophysical Research*, **92**, 6805-6824.
9. Fily, M. and Rothrock, D.A. (1986), "Extracting Sea Ice Data from Satellite SAR Imagery," *IEEE Transactions on Geoscience and Remote Sensing*, **24**, 849-854.
10. Fily, M. and Rothrock, D.A. (1987), "Sea Ice Tracking by Nested Correlations," *IEEE Transactions on Geoscience and Remote Sensing*, **25**, 570-580.
11. German, S., and German, D. (1984), "Stochastic relaxation, Gibbs distributions and the Bayesian restoration of images," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **6**, 721-741.
12. Hastie, T., and Stuetzle, W. (1989), "Principal Curves," *Journal of the American Statistical Association*, **84**, 502-516.
13. Murtagh, F., and Raftery, A.E. (1984), "Fitting Straight Lines to Point Patterns," *Pattern Recognition*, **17**, 479-483.
14. Ninnis, R. M., Emery, W. J., and Collins, M.J. (1986), "Automated Extraction of Pack Ice Motion From Advanced Very High Resolution Radiometer Imagery," *Journal of Geophysical Research*, **91**, 10725-10734.
15. Rosenfeld, A., and Kak, A.C. (1982) *Digital Picture Processing* (2 vols.), New York: Academic Press.
16. Rothrock, D., and Thorndike, A. (1984), "Measuring the Sea Ice Floe Size Distribution," *Journal of Geophysical Research*, **89**, 6477-6486.
17. Serra, J. (1982) *Image Analysis and Mathematical Morphology*, New York: Academic Press.
18. Vesecky, J.F., Samadani, R., Smith, M.P., Daida, J.M. and Bracewell, R.N. (1988), "Observation of Sea-Ice Dynamics Using Synthetic Aperture Radar Images: Automated Analysis," *IEEE Transactions on Geoscience and Remote Sensing*, **26**, 38-48.
19. Ward, J.H. (1963), "Hierarchical Grouping to Optimize an Objective Function," *Journal of the American Statistical Association*, **58**, 236-244.

Profiles in Science



Murray Rosenblatt

Murray Rosenblatt has been a Professor of Mathematics at the University of California, San Diego in La Jolla since 1964 and a principal investigator of the Office of Naval Research for many years. He received his undergraduate education at the City College of New York and his graduate education at Cornell University. His interests have generally been in the areas of probability theory and mathematical statistics as well as allied problems in mathematical analysis. Some of his early research centered about the domain of validity of the central limit theorem (asymptotic normality) for dependent random sequences (narrow band pass filtering) as well as the design of spectral estimates. He has also been concerned with probability density and regression estimation

from a nonparametric point of view.

More recently he has considered questions related to dependent random sequences with long range dependence and the application of methods of spectral estimation (of higher order) in the analysis of turbulent velocity fields. Also, non-parametric and parametric methods as applied to nonGaussian possible nonminimum, phase random processes or fields are topics of current interest. Such methods are relevant in resolving a class of deconvolution problems arising, for example, in geophysical analysis. Dr. Rosenblatt was a Guggenheim fellow in the years 1966 and 1972. He is a member of the International Statistical Institute and the National Academy of Sciences.

Neural Networks: A Computer Science Perspective

Dr. Alan L. Meyrowitz
Navy Center for Applied Research in Artificial Intelligence

Introduction

The pace of basic research in neural networks has been accelerating for several years, with projects at academic and industrial laboratories seeking to answer fundamental questions in mathematically modeling neural processes, explicating the biology of organic neural networks, and implementing artificial neural networks in VLSI electronics. As mathematicians, biologists, and electronics engineers have been drawn to the study of neural networks, so have computer scientists. However, in contrast to the particular motivations for research in the other disciplines, computer scientists are fundamentally driven by a desire to understand what can be computed, and with what efficiency. For example, once it has been determined that some computational task can in fact be accomplished by a network, it is of interest to determine what lower bounds may exist on the time required for computation, or on the number of processors which must be made available. What has emerged to date has been the understanding that artificial neural networks can fill an important niche in implementing the computation required for some special classes of problems at the lowest levels of perception and memory; they do so by effectively functioning as an associative memory. This entails

the learning of a mapping between an input space and an output space, such that the mapping generalizes properly when applied to new inputs (Poggio 1990). Artificial neural networks have consequently been studied intensively for their applications to pattern recognition and signal processing where such mappings are of central importance, but they fall far short of providing any sort of "silver bullet" for dealing with computational problems generally.

This research arena is better characterized by the term *connectionism*, rather than by *neural networks*. Connectionism quite properly has come to connote the interest in understanding what can be accomplished computationally by massively parallel interconnections among processing units. On the other hand, use of the phrase neural networks exaggerates the biological inspiration for work in progress, at least within the computer science research community. Collaboration with the life sciences has been fruitful, of course, and for example has had precedent in the progress made in machine learning; in many ways that progress has been built on the coupling of insights from both cognitive psychology and computer science. Still, powerful machine learning techniques have been discovered that seem to have no psychological plausibility, and similarly (notwithstanding a lack of biologi-

cal motivation) a broad spectrum of possible connectionist approaches are being considered that have potential utility in the realm of artificial systems.

There are possibilities beyond the biological

The range of possibilities is immediately apparent when one considers the variables that make up the general connectionist model.

As a processing unit functions within a network, it receives as numeric input messages from other units, applies a function to those messages to update its internal activation level, and generates an output available to other units to which it is connected. The connecting links have numeric weights associated with them, used to multiply the numeric input of the unit to which the link is attached. Computer scientists have a great deal of freedom in selecting for study particular functions which can be presumed to determine new activation levels within a unit, on the basis of current activation level and input received, or to determine changes in the weighting of links on the basis of feedback to the network regarding its global behavior (see Figure 1). There is in fact a growing corpus of important connectionist literature dealing with models not purporting to be physiologically correct, yet contributing to a deeper understanding of the role which massive parallelism can play in the design of computer systems. To cite just one example, the fundamentally important capability to do pattern matching among strings of symbols has been demonstrated to be a property that emerges from collections of neuron-like elements (Touretzky 1985). In light of such results, it is quite possible that computer science research will

be fruitful in leading to artificial systems with brain-like capabilities long before the nature of any corresponding biological implementation is clarified. Indeed, following research directions intended to merely mimic the brain must be regarded as unnecessarily constraining the possible technology. Approaches to knowledge representation and automated reasoning, independent of biological studies, could very well result in computer systems surpassing human performance in learning, planning, general problem solving, and other intellectual skills.

It might be argued that the efforts of computer science ought to be focused as much as possible on the apparently biologically correct models, given the impressive capabilities of the brain and the brain's existence as proof of what can ultimately be expected of artificial systems which adhere to those models. This argument is flawed for a number of reasons, including the fact that we cannot yet point to any models which capture the complexity of the brain's neural architecture in any but the most superficial way; there is as yet no deep understanding of what accounts for the information processing in a single neuron, let alone what accounts for computation among neurons in their massively parallel organization. Consequently, we are not ready to predict what particular network structures or processing paradigm may emerge as most crucial to the brain functions that we would have our computer systems mimic (Schwartz 1988; Rumelhart 1990).

Furthermore, it is also conceivable that once the mechanisms underlying the brain's capabilities are revealed, either in low-level sensory data processing or in high-level cognition, we will discover that they do not map well into corresponding mechanisms suited to the silicon and copper of computers. (See Figure 2)

Figure 1

The position of a simple network: Functions f , g , h are usually weighted linear combinations of the inputs. Computed values above a set threshold will determine each node's output.





Figure 2

Typically, output is a nonlinear function of the threshold value computed at each node.



Symbols remain a viable alternative to patterns of connections

Only a small portion of the computer science research community has embraced the neural approach as a pathway toward artificial intelligence. The notion that intelligence can be derived from the changing pattern of connections among processing units stands in contrast to the mainstream hypothesis articulated early on by Allen Newell (Newell et al. 1972) that intelligence is fundamentally the manipulation of symbols. It still remains a legitimate point of view to suggest that within the brain, as we move away from the processing of sensory data toward the higher level cognition, we will see that processing is functionally equivalent to the manipulation of symbols. Computers, being extremely adept at symbol manipulation, can take that level of processing as their starting point and can potentially be made to at least duplicate human cognitive skills. From this perspective, the research of cognitive psychologists is more relevant to computer science interests than is the research of neurophysiologists. It is worth noting that Newell has amassed a great deal of evidence in support of symbol manipulation as the basis for intelligence in cognitive systems, both artificial and human. His experiments with Soar (Newell 1990), an experimental architecture for general intelligence, have demonstrated that many varieties of human cognition, ranging from skill in game playing to theorem proving and more general problem solving, can be at least closely approximated by system designs which assume symbol manipulation is in the form of processing situation/action (or production) rules of inference.

Hybrid approaches overcome neural deficiencies

Exploiting network processing alone and seeking to adhere to biological plausibility has resulted in systems severely limited in their capabilities, as one may observe in the arena of research in natural language understanding which has been slow to demonstrate parsing beyond fixed-length sentences (Waltz, et al 1985; Cottrell 1985; Hanson et al. 1987). In the same vein, network implementations of situation/action pairs have the potential to mimic the deductive capability of the sets of inference rules in expert systems, but provide no basis for generating the human-understandable explanations of the deduction; such explanations are crucial to the debugging of new systems and to the confident use of operational systems. Seeking to accelerate progress by dispensing with the constraints imposed by biological plausibility can take a variety of forms, such as exploring the conjunction of separate but coordinated neural and symbolic processors in so-called hybrid systems (necessitating the addressing of issues in controlling their activities and in supporting communication between them), or simply giving free reign to the

building of networks with nodes and links having an explicit meaningful association with specific objects and relationships in the world. The latter approach opens the door to representing particular inference rules as specific linkages among designated nodes, where the number of nodes is dictated by the number of rule clauses to be represented and weights and thresholds at a node are chosen to faithfully represent the required logical conjunction of clauses. An explanation of network behavior can be obtained by interpreting node activation in terms of the associated rules (Shavlik 1991).

We thus see the potential for expert system functions to be implemented by an approach which is neurally-oriented, but clearly a significant degree of thoughtful design must go into the effort in a way that is quite analogous to the intellectual effort which programmers must bring to the task of software engineering. Unfortunately, the occurrence of somewhat glib statements in the literature and other statements quoted out of context have generated a perception that artificial neural networks are obviating the need for programming and, by implication, the need for research in software engineering. For example (Reilly 1990),

"Since neural networks learn, they differ from the usual artificial intelligence systems in that the solution of real-world problems requires less of the expensive and elaborate programming and knowledge engineering required for such products as rule-based expert systems."

Such a statement can mislead one to assume that neural networks, by virtue of their learning capability, have actually duplicated the functionality of expert systems and other complex software. That is clearly not the case, and it remains an extreme extrapolation of empirical evidence to date to assume that it will ever be the case.

When hybrid systems can be built that have essentially distinct though at least loosely coupled neural and symbolic components, it is reasonable to expect choices in system design to be available for exploiting the best features of both kinds of processing. It would be wrong, however, to assume that requirements for high speed or robustness in computation - attributes often claimed for networks - will strictly dictate the division of labor in a hybrid system, with the slower, well-defined tasks having accurate and complete input available being reserved for the symbolic component. Symbol manipulation systems can exhibit rapid retrieval (measured in microseconds) of complex information by means of content-addressable memory, and appropriate logic can make programs robust as software bugs or unexpected input are encountered (de Callatay 1986). It is thus the case that, by a rather straight forward extrapolation of current technological approaches, we can foresee the claimed advantages of neural networks being realized by symbol manipulation systems.

Original from

UNIVERSITY OF ILLINOIS AT URBANA-CHAMPAIGN 15

In contrast, there appear to exist grand challenges to achieving with neural networks the computational capabilities afforded by symbol manipulation, and that is the case whether one considers artificial intelligence or more conventional computation. The latter, for instance, is very much dependent on the ability to accomplish variable binding, i. e., assigning values to global variables either dynamically during the course of information processing or permanently in memory, as might be required to store possible instances of a general schema. It has been observed (Waltz et al. 1988) that variable binding has a serial character and there has not yet emerged a satisfactory connectionist approach to it.

We can exploit results of machine learning research

With regard to artificially intelligent systems, the steady progress which has been made in the special topic of machine learning has demonstrated many times over the importance of using knowledge already acquired in order to facilitate the acquisition of new knowledge. We can more generally consider prior experience as having the effect of altering the structure of a system, ideally in a beneficial way as part of a maturing process for dealing with a complex, unpredictable environment. There is a strong argument to be made, therefore, for imposing some initial structure on artificial neural networks as we build them; certainly, the learning efficiency of networks is very much dependent upon the initial choice of weights on the links between the units (Pollack 1989). Moreover, those researchers who aspire to follow a biological model are bound to credit DNA encoding with dictating how neural tissue develops over time, imposing a network structure which reflects compiled evolutionary experience. With the goal in mind of exploring the broad range of possibilities for designing computer systems, it makes sense to begin with the structures that have already been demonstrated to have merit, such as semantic networks (where nodes are individually associated with objects or concepts and links with relationships among them). For instance, semantic networks can be augmented by adjustable weights on links that in traditional representations would simply denote fixed relationships. This approach to automating high-level cognition constitutes what has come to be called structured connectionism, and offers perhaps the best prospect for achieving an advantageous, intimate merging of neural and more conventional techniques.

Advances have been in recognition, not understanding

Progress in the VLSI implementation of artificial neural networks has been most impressive in the work of Carver

Mead and his colleagues (Lazzaro et al. 1990), in particular in building a 220,000 transistor chip modeling the time-domain processing that accounts for auditory localization in the barn owl. Still, this processing of sensory data provides only the initial information which would be the basis for an intelligent system's representation of its environment. It is the processing most removed from, and preliminary to, the high-level cognitive processing that is of central interest to researchers in intelligent systems.

In fact, it is generally the case that what artificial neural networks can provide in the way of computation, by closely adhering to the best understood examples of biological neural processing, is efficient pattern recognition. However, this is not viewed as dramatically impacting computer scientists pursuing artificial intelligence. That community came to realize quite some time ago (in attempts at natural language processing and the machine translation of language) that recognition is distinct from understanding, and that in many ways recognition is a simpler problem than understanding. Recognizing an object as a submarine is one thing; quite another is understanding the implications of its appearance in the context of a given situation.

Lack of an appreciation for the distinction between recognition and understanding can lead one to conclude, incorrectly, that the discrimination of subclasses of objects in a population, where those subclasses might be missed by a conventional symbolic induction algorithm, represents a triumph of the neural approach over conventional machine learning techniques. Rather, the situation is that symbolic induction has a different goal, and arguably a more ambitious goal in yielding not just a partition of sensed data into subclasses but in addition yielding conceptually meaningful descriptions of those subclasses. The descriptions, which can take the form of classification rules, allow an insightful interpretation of object classes by humans and can support higher level reasoning about the classes by machines. Significant progress (Michalski et al. 1990) toward the generation of such descriptions has been made by the machine learning research community, and has been important in building the scientific basis for automating understanding. Trainable classifiers and self-organizing classification networks do not provide those descriptions; nor for that matter do conventional statistical clustering techniques (Michalski 1982).

Conclusion

We have pointed out a number of facets of the state-of-understanding regarding neural networks which should clarify the observations already made, that there is minimal enthusiasm for neural processing as a basis for artificially intelligent systems, and that neural processing does not constitute a revolutionary advance for computer science generally. Indeed, we noted at the outset the primary interest in understanding

Original from

UNIVERSITY OF ILLINOIS AT
URBANA-CHAMPAIGN

what can be computed, and in understanding the theoretical bounds on computational efficiency. Artificial neural networks have thus far not broadened the scope of functions or processes known to be computable. Tasks such as pattern recognition for which they are suited can be accomplished, albeit less efficiently, by conventional symbolic, algorithmic means, making neural networks the technology of choice for those tasks at this time. However, there is as yet no theoretical basis for saying that the advantages of neural networks for those tasks cannot be duplicated or surpassed by the evolution and occasional revolution which is occurring in conventional techniques.

To the extent that there is an indisposition toward neural networks in the artificial intelligence research community, it is not due so much to an intuition about their potential inadequacy, but rather can be attributed to the genuine excitement about the progress being made by symbol processing approaches. For example, it was not that long ago that machine learning was viewed as a grand challenge; today numerous techniques vie for attention and comparative study. Neural computation may contribute to overcoming the grand challenges that remain, but there is little at the moment to suggest the contribution will be significant. For example, there is considerable progress being made in basic research, building the science base for the future technology of autonomous, intelligent robots capable of functioning in dynamic and unpredictable environments. Such a robot may very well be a hybrid system, with a neural component providing the initial filtering and processing of sensory data. The symbol processing component is envisioned as providing the cognitive capability for planning, problem solving, and understanding. In light of the momentum that has been achieved in recent years toward making the latter component a reality, there is a natural and justified reluctance to divert attention and resources toward the neural arena without more evidence of its potential.

The latter subjective view is buttressed by the still small but growing body of results from comparative studies of neural algorithms and conventional machine learning algorithms, results that have been far from anything that might dramatically shift the agenda of the research community in artificial intelligence. For example, experiments have been conducted comparing a prominent symbolic learning algorithm (ID3) with back-propagation neural learning algorithms, using several large real-world data sets; back-propagation performed about the same as ID3 in terms of classification correctness on new examples, but took much longer to train (Shavlik et al. 1989). Other studies applied statistical pattern recognition, neural networks, and machine learning to four real-world data sets, with detailed attention being given to the analysis of performance of the neural networks using back-propagation; overall, the networks were judged not to be the best classifiers, they consumed enormous amounts of cpu time, and the results suggested that improved performance would have to await

further research progress in network training and representation (Weiss et al. 1989).

Moreover, these empirical insights into performance deficiencies in artificial neural networks may reflect theoretical limitations on what such networks can accomplish. It is already known that a number of network design and learning problems are NP-complete, which is to say that only small instances of the problems will admit exact solutions and that for other instances, even with the expenditure of large amounts of computer time, only approximate solutions can be found. For example, there is evidence that training neural networks is intrinsically difficult computationally, for certain tasks on the basis that NP-completeness has been proved for problems in determining the existence of threshold functions, and in assigning weights to links in an arbitrary network (Judd 1987; Blum et al. 1990). The implication is that severe constraints may exist on the theoretic ability of artificial neural networks to be efficient in accomplishing tasks where resources are limited; the fact that the back-propagation algorithm has been observed to run slowly may be a reflection of those constraints. It remains a worthwhile endeavor to determine the specific learning and problem-solving tasks for which, as a practical matter, the availability of those resources is sufficient to allow the neural approach to be used to advantage.

Biography

Dr. Alan L. Meyrowitz is Director of the Navy Center for Applied Research in Artificial Intelligence at the Naval Research Laboratory. Formerly he was Manager of the Intelligent Systems Program in the Computer Science Division of the Office of Naval Research. His research interests are in artificial intelligence and robotics, with a focus on machine learning and autonomous systems.

References

1. Blum, A. et al. Training a 3-Node Neural Network is NP-Complete, Technical Report, MIT Laboratory for Computer Science, Cambridge 1990.
2. Cottrell, G. W. Connectionist Parsing, Proceedings of the Seventh Annual Conference of the Cognitive Science Society, Irvine 1985.
3. de Callatay, *Natural and Artificial Intelligence*, North-Holland, Amsterdam 1986.
4. Hanson, S. J. et al. A Connectionist Network That Learns Natural Language Grammar, Proceedings of the Ninth Conference of the Cognitive Science Society, Seattle 1987.
5. Judd, J. S. Complexity of Connectionist Learning With Various Node Functions, University of Massachusetts Technical Report COINS 87-60, July 1987.

6. Lazzaro, J. et al. A Silicon Model of Auditory Localization, in *An Introduction to Neural and Electronic Networks* (ed. S. F. Zornetzer, J. L. Davis, C. Lau), Academic Press 1990.
7. Michalski, R. S. et al. *Machine Learning*, volume 3, Tioga, Palo Alto 1990.
8. Michalski, R. S. et al. Learning From Observation: Conceptual Clustering, in *Machine Learning*, volume 1, Tioga, Palo Alto 1982.
9. Newell, A. et al. *Human Problem Solving*, Prentice-Hall, Englewood Cliffs 1972.
10. Newell, A. *Unified Theories of Cognition*, Harvard University Press, Cambridge 1990.
11. Poggio, T. et al. Networks for Approximation and Learning, Proceedings of the IEEE, vol. 78, no. 9, September 1990.
12. Pollack, J. B. Connectionism: Past, Present, and Future, Artificial Intelligence Review, vol. 3, 1989.
13. Reilly, D. L. et al. An Overview of Neural Networks, in *An Introduction to Neural and Electronic Networks* (ed. S. F. Zornetzer, J. L. Davis, C. Lau), Academic Press, 1990.
14. Rumelhart, D. E. Brain Style Computation, in *An Introduction to Neural and Electronic Networks* (ed. S. F. Zornetzer, J. L. Davis, C. Lau), Academic Press 1990.
15. Schwartz, J. T. The New Connectionism, DAEDALUS, winter 1988.
16. Shavlik, J. W. Directed Propagation of Training Signals through Knowledge-Based Neural Networks, Proceedings of the Twelfth International Joint Conference on Artificial Intelligence, 1991.
17. Shavlik, J. W. et al. Symbolic and Neural Learning Algorithms: An Experimental Comparison, University of Wisconsin Technical Report CSTR #857, June 1989.
18. Touretzky, D. S. et al. Symbols Among the Neurons: Details of a Connectionist Inference Architecture, Proceedings of the International Joint Conference on Artificial Intelligence, 1985.
19. Waltz, D. et al. *Connectionist Models and Their Implications*, Ablex, Norwood 1988.
20. Waltz, D. et al. "Massively Parallel Parsing," *Cognitive Science*, 9, pp 51-74, 1985.
21. Weiss, S. M. et al. An Empirical Comparison of Pattern Recognition, Neural Nets, and Machine Learning Classification Methods, Proceedings of the International Joint Conference in Artificial Intelligence, Detroit 1989.

Intelligent Tutoring Technology: The Emergence of Individualized Instruction

*Bruce W. Hamill**
Johns Hopkins University
Applied Physics Laboratory

Imagine learning a new computing language from an intelligent tutoring system that presents you new lessons to learn and problems to work on by taking into consideration your individual level of progress and areas of difficulty, as determined dynamically by your performance on the problems it presents to you, and your personal learning strategy, which it infers by watching you work on problems and matching your approach to one of many that it knows about. Sound like science fiction? It is not. Such a system has been used to teach the introductory LISP programming course at Carnegie-Mellon University, and its students have gotten better grades than students who do programming exercises in lab sessions with human instructors!

Individualized computer-based instruction of this kind is a revolutionary advancement in instructional technology. This impressive new technology is emerging from very recent research in cognitive science and artificial intelligence. Building on some two decades of research in computer-assisted instruction, intelligent tutoring systems are being developed to train individuals to perform a range of complex academic and real-world problem-solving tasks (See note 1). Among the subject areas being addressed are geometry, economics, computer programming, medical diagnosis, electronic troubleshooting, and system maintenance. While many of these intelligent tutoring systems remain vehicles for research on human skill learning and cognitive performance, some have been implemented as operational systems and others are un-

dergoing transitions from research to operational environments.

Such systems have the potential for dramatically altering the manner in which students and trainees learn new subjects because they are designed to adapt to and take advantage of the different ways in which individuals learn. They are tolerant of mistakes, and they can bring such mistakes to the student's attention, diagnose them, and use those diagnoses to guide the student to a correct solution, as would a human tutor.

Research Foundations

Intelligent tutoring technology rests on the foundation of scientific investigations of human cognition by researchers in cognitive science and its contributing disciplines, including cognitive psychology, education, artificial intelligence, and computer science. In these investigations, the focus is on understanding how individuals learn and remember new material, solve problems, and acquire skill and expertise: how they organize and approach new material, what knowledge they bring to bear, how that knowledge is organized, what learning and problem-solving strategies they use, what kinds of errors they make along the way, how they recover from errors, and even the nature of misconceptions that result from different kinds of uncorrected learning and problem-solving failures.

Over twenty years ago, the Office of Naval Research (ONR), recognizing the potential value of research on com-


puter-assisted instruction, began making foresighted research investments which made possible the advances we are seeing today¹. (See note 1) Among the earliest products of these research investments were the seminal studies of Jaime Carbonell and Allan Collins and their colleagues at Bolt Beranek and Newman, Inc., a private research laboratory. In an influential 1970 paper, Carbonell introduced the concept of "information-structure-oriented" computer-assisted instruction². The information structure consisted of a semantic network of concepts and their relationships, a structure that would permit flexible "mixed-initiative dialogues" between the student and the tutoring system wherein the system could respond to student questions, including unexpected questions, and could ask questions to lead the student toward mastery of the subject matter, all by traversing its semantic network of information about the subject. This was in contrast to the earlier concept of "frame-oriented" computer-assisted instruction in which detailed instructional information was stored in fixed frame structures supporting answers to expected student questions

about the subject matter, together with predetermined branching decisions to the next units of instruction based on the student's answers to system-generated questions. Carbonell's new concept was exemplified in his system called SCHOLAR, which tutored students in Latin American geography. It was an important first step toward the development of intelligent tutoring systems.

To advance their understanding of the processes involved in reasoning and in tutorial dialogues, Collins and his colleagues studied the Socratic method of instruction, in which the tutor leads the student to principles of knowledge not by expository statements but by a process of asking the student questions and encouraging him to use principles of reasoning to discover and validate the targeted principles of knowledge for himself³. Over the course of several productive years of such research, they developed and evaluated computer-assisted instruction systems and strategies for tutoring causal knowledge and reasoning, and they developed a computational theory of tutoring. Current tutoring theory and tutoring

Figure 1

Lt. Bill Marriott at a simulation-based intelligent tutoring system for marine steam-propulsion power plant developed by T. Govindaraj and his colleagues at Georgia Institute of Technology.



strategies derive in large part from their analyses of human tutorial dialogues, which were conducted in terms of how such dialogues could be implemented in computers.

It is useful from one theoretical perspective to view the goal of instruction as the communication of knowledge from the instructor to the student⁴. In order for this to be accomplished effectively on an individual basis, there are several requirements for matching the instructor's knowledge and instructional strategies to the student's knowledge and learning strategies. The instructor's attempts to convey new knowledge to the student will be thwarted if the student's knowledge base lacks the fundamental structure and content for assimilating the new knowledge (for example, it is difficult to learn trigonometry without adequate grounding in algebra and geometry). Similarly, a mismatch between the instructor's teaching strategies and the student's learning strategies may lead to failure in the communication of knowledge (for example, when two different paths lead to the same goal, and the instructor is taking one path while the student is trying to take another).

William Clancey of the Institute for Research on Learning has addressed these problems by designing instructional systems to teach students how to perform medical diagnoses. He started by building a tutoring system (GUIDON)⁵ to work with the knowledge base of MYCIN, an artificial intelligence-based expert consultation system designed for diagnosing infectious diseases⁶. In the course of developing GUIDON and analyzing expert human teachers' instructional dialogues with students, Clancey observed structural and procedural regularities relating to medical diagnosis and to instruction that had not previously been taken into account. These human characteristics, such as generating and evaluating alternative hypotheses, focusing on selected hypotheses, and using strategic knowledge to guide search through the knowledge base, led him to successive reformulations of both the expert consultation system and the tutoring system. Clancey's redesigning of these systems reflects his important discovery of the difference between expert knowledge of a subject and the knowledge required to teach that subject. In addition to developing models of the subject matter and the instructional procedure, he considered the related problem of the student model, that is, the system's qualitative model of the student's knowledge of the subject matter and how it changes in the course of instruction⁷.

Qualitative student models of this kind require deep understanding of the processes of human learning. Cognitive scientists have been studying individual differences in learning and in other cognitive task performance by analyzing how students approach and perform tasks during learning and problem solving. Detailed analyses of such performance data permit computational models of student learning and problem-solving performance to be constructed in the form of computer programs. When these programs are run as simulations of students performing similar tasks, they produce the

kinds of performance, both correct and errorful, that are characteristic of real students performing such tasks.

From Theory to Applications

Cognitive science research and intelligent tutoring system development have been a focus of the work of John Anderson of Carnegie-Mellon University. His work is based on his theory of cognitive skill acquisition (how people develop expertise) called ACT*, which provides a structure for describing the organization of knowledge and its use in cognitive tasks like learning and problem solving⁸. Using his theory as a foundation, he has built intelligent computer systems to tutor students in several subjects, including geometry⁹ and the computer programming language LISP¹⁰; some results of using the LISP Tutor were mentioned at the beginning of this article.

These intelligent tutoring systems serve as laboratories for studying human learning and problem solving. In operation, the tutoring systems collect data on students' keyboard responses, thereby providing detailed records of their choices through the entire course of each tutorial session. Experiments can be conducted by manipulating instructional conditions, and results of data analyses can guide principled changes to the structure and operation of the tutoring systems. As his understanding of processes of learning improves through such experimentation, Anderson makes appropriate modifications to his learning theory and records them in computational form in his intelligent tutoring systems, which thus become tangible applications of his theoretical model of learning and problem solving.

A different approach to intelligent tutoring, which is being taken by Robert Glaser and his colleagues at the University of Pittsburgh's Learning Research and Development Center, is to design "intelligent discovery worlds" in which students can explore new concepts in a subject area, such as economics¹¹. (See note 2) The discovery world is a computer simulation program that provides students an interactive environment in which to engage in inductive or discovery learning. In an economics microworld, known as Smithtown, on-line tools are provided to support the student's informal exploration of economic concepts by manipulating values of variables such as prices and quantities of goods, conducting informal experiments to determine relationships among those variables, and inducing general laws, such as the law of supply and demand, from the results of their experiments. Microworlds permit students to proceed in their own directions, using their own individual learning strategies to explore new subjects. The history records of these explorations enable the system (and human instructors) to identify specific weaknesses in a student's approach to problems, as well as to confirm successful accomplishment of tasks. Future plans involve using what is learned about the strategies of successful students to develop guided discovery aids to support explorations of the microworld.

Other approaches to the design of intelligent tutoring systems are also being taken. James Hollan and Edwin Hutchins and their colleagues, then at the Navy Personnel Research and Development Center and the Institute for Cognitive Science at the University of California, San Diego, developed an interactive simulation-based instructional system for steam-propulsion plants of the kind that power Navy ships. This system, known as STEAMER, is based upon a detailed mathematical model of a complex steam-propulsion power plant, with all of its components and their interactions¹². It is intended to provide practice on a realistic simulation of an operational system; it has been used successfully for training at the Great Lakes Naval Training Center. The workings of this model may be observed at different levels of detail through a highly informative color graphics interface that supports detailed displays of systemic effects of any conditions that can be simulated by changing parameters in the underlying mathematical model. This permits students to observe such things as the distribution and flow of the main engine lube oil system or the feedwater system, changes in dial and gauge readings resulting from changes in settings of pumps, valves, and other components, and (possibly catastrophic) effects of incorrect settings of valves, switches, and other system controls. STEAMER's designers consider the detailed graphical interface to be a representational system designed for communication between the system and the student, and they deem it essential to support the interface with detailed understanding of the cognitive task that the system is attempting to support¹³.

Although STEAMER was never fully developed as an intelligent instructional system, it influenced other related work. A simulation-based intelligent tutoring system for marine steam-propulsion power plants has been developed by T. Govindaraj and his colleagues at Georgia Institute of Technology¹⁴. Instead of building a detailed quantitative model, as was done in STEAMER, Govindaraj used a qualitative approximation methodology in which states of his simulation system (Turbinia) are represented as qualitative functional descriptions of a hierarchy of subsystems, components, and primitives (the basic functional units) (Figure 1). Structures of the primitives are based on approximate functional equivalents, rather than exact values, of the appropriate differential and algebraic equations of system dynamics, and each primitive has a set of parameter values associated with the particular component that it represents. Changing states of the system evolve through updates to individual components which are then propagated to successor components, maintaining temporal fidelity in the sequence of events. Perturbations of the system are represented as deviations of numerical state values from their nominal values. These numerical values represent system states only qualitatively, since they are derived from approximate system equations. They are transformed into qualitative descriptions, such as "pressure low" or "level high" before they are presented to trainees. This approximate, qualitative represen-

tation of system states enhances cognitive compatibility between Turbinia and trainees to the extent that such representations are similar to state descriptions used by trainees.

Turbinia is linked to an intelligent tutoring system (Vyasa) to provide instruction in diagnostic troubleshooting. Vyasa contains knowledge, organized at several levels of granularity, about Turbinia's structure, function, and behavior, and about troubleshooting. It also contains knowledge of instructional strategies and keeps track of trainee progress, errors, and possible misconceptions, offering help as needed and when requested by the trainee.

Turbinia-Vyasa is implemented in Common LISP with object-oriented extensions and runs on Apple Macintosh II machines, making use of the graphics, icons, and mouse-based menu-selection capabilities in a schematic interface. An experimental evaluation of Turbinia-Vyasa, in which Navy ROTC students served as subjects, compared troubleshooting performance on a set of problems in the Turbinia simulator by groups undergoing either unaided training on Turbinia or training on Turbinia that was coupled with the use of the Vyasa tutor. Results indicated that Vyasa helped trainees develop good troubleshooting strategies, including formation of plausible failure hypotheses based on observed symptoms and systematic elimination of such failure hypotheses by the conduct of appropriate diagnostic tests; those trained without the Vyasa tutor did not develop good troubleshooting strategies, instead relying heavily on guessing. Results also revealed differential effects of Vyasa's passive (provide help upon request and without intervention) and active (with intervention) modes of providing instruction and help to trainees: Those trained with the active tutor made fewer guesses and fewer premature diagnoses than those trained with the passive tutor, but they also solved fewer problems in the time allotted, and a few of them became dependent on Vyasa for help in solving problems. This suggests that trainees are not all equally receptive to particular tutoring strategies, and that consideration must be given in the design of intelligent tutoring systems to individual differences in abilities and preferences.

Douglas Towne and his colleagues at the Behavioral Technology Laboratories of the University of Southern California recently completed development of their Intelligent Maintenance Training System (IMTS), which employs a generalized approach to training operators to diagnose problems in hardware devices, together with editing tools for rapidly producing a specialized training system for each device. This system evolved from several successive systems that were built over the course of twenty years, each taking advantage of new developments in computer hardware and software technology and improved understanding of human problem solving, learning, and instruction processes¹⁵. IMTS comprises a set of models needed to support effective instruction in troubleshooting: an instructional model, which is a set of functions that manage the instruction and problems provided

to individual students on the basis of demonstrated progress and difficulties; a model of the student's mastery of domain information that the instructional model constructs dynamically from results of the student's success in solving problems; an expert troubleshooter model, which is a set of device-independent processes that evaluate the student's diagnostic approach, assist the student as necessary, and demonstrate preferred diagnostic techniques; and a model of the device on which the training is to be conducted. IMTS also has authoring tools that allow experts to create equipment-specific simulations that will be responsive to system configuration alterations made by students as they work on troubleshooting problems.

Conclusion

Intelligent computer-assisted instruction technology is an increasingly important part of civilian and military education and training. The tutoring systems described here share one crucial characteristic: They are direct products of research on processes involved in human learning and instruction. They are explicitly designed to recognize and to take advantage of human cognitive performance abilities and limitations. And they have shown that this is what it takes to build truly intelligent tutoring systems.

Further Reading

A number of recent publications present discussions of theory and empirical results that can lead the interested reader to the current state of the art of intelligent tutoring systems research and development. These include the proceedings of the 1988 international conference on intelligent tutoring systems¹⁶, edited collections of seminal theoretical and empirical papers¹⁷⁻²⁰, and textbooks^{4,21}. Hamill²² offers some directions for future research.

Biography

Bruce W. Hamill earned his Ph.D. at Johns Hopkins University, where he studied psycholinguistics and cognitive psychology. Since 1979 he has been a research psychologist at the Johns Hopkins University Applied Physics Laboratory. He has worked on a variety of naval problems, including visual perception of symbolic and acoustic data, intelligibility of helium speech, tactical decision making, cognitive engineering or rule-based decision support systems, and cognitive science issues in instruction and training.

Notes

* This article was prepared while the author served as a Scientific Officer in the Cognitive Science Program of the

Office of Naval Research (ONR) on an Intergovernmental Personnel Mobilization Act (IPA) assignment. I am grateful to Susan Chiptman, Program Manager for Cognitive and Decision Sciences, for helpful comments on an earlier draft of this article.

Please send requests for reprints to Bruce W. Hamill, Johns Hopkins University Applied Physics Laboratory, Johns Hopkins Road, Laurel, MD 20723-6099.

1. ONR was willing to explore the possibility of intelligent computer-assisted instruction systems despite the prevailing impracticality of the necessary computing facilities. Most of the investigators whose efforts are described here have received ONR support at some time during the past two decades. ONR also provided state-of-the-art computing systems to several university laboratories through a Defense University Research Instrumentation Program (DURIP) grant to the University of Pittsburgh's Learning Research and Development Center (LRDC). Under this grant, LRDC researchers developed software tools for use in intelligent tutoring projects in civilian and Department of Defense laboratories.

2. Glaser's work was supported by the Department of Education, Office of Educational Research and Improvement; ONR funded the acquisition of computer equipment for the work under the DURIP grant to LRDC.

References

1. Office of Naval Research. (July 1969). *ONR Psychological Sciences Division Programs* (Report 450-1, other reports in series). Arlington, VA: Author.
2. Carbonell, J. R. (1970). AI in CAI: An artificial intelligence approach to computer-assisted instruction. *IEEE Transactions on Man-Machine Systems*, 11(4), 190-202.
3. Collins, A., Warnock, E. H., Aiello, N., & Miller, M. L. (1975). Reasoning from incomplete knowledge. In D. G. Bobrow & A. Collins (Eds.), *Representation and understanding: Studies in cognitive science*. New York: Academic Press.
4. Wenger, E. (1987). *Artificial intelligence and tutoring systems: Computational and cognitive approaches to the communication of knowledge*. Los Altos, CA: Morgan Kaufmann.
5. Clancey, W. J. (August 1986). From GUIDON to NEO-MYCIN and HERACLES in twenty short lessons: ONR final report 1979-1985. *AI Magazine*, 40-187. Vol 7: (3) 40-60, 187.
6. Shortliffe, E. H. (1976). *Computer-based medical consultations: MYCIN*. New York: American Elsevier.
7. Clancey, W. J. (1986). Qualitative student models. *Annual review of computer science*, Palo Alto, CA: Annual Reviews, Inc., 381-450.
8. Anderson, J. R. (1983). *The architecture of cognition*. Cambridge, MA: Harvard University Press.

Original from

UNIVERSITY OF ILLINOIS AT 23
URBANA-CHAMPAIGN

9. Anderson, J. R., Boyle, C. F., & Reiser, B. J. (1985). Intelligent tutoring systems. *Science*, **228**, 456-462.
10. Anderson, J. R., & Reiser, B. J. (1985). The LISP tutor. *BYTE*, **10**(4), 159-175.
11. Shute, V., Glaser, R., & Raghavan, K. (1989). Inference and discovery in an exploratory laboratory. In P. L. Ackerman, R. J. Sternberg, & R. Glaser (Eds.), *Learning and individual differences: Advances in theory and research*. San Francisco: Freeman.
12. Hollan, J. D., Hutchins, E. L., & Weitzman, L. (1984). STEAMER: An interactive inspectable simulation-based training system. *AI Magazine*, **5**(2), 15-27.
13. Hollan, J. D., Hutchins, E. L., McCandless, T. P., Rosenstein, M., & Weitzman, L. (1987). Graphic interfaces for simulation. In W. B. Rouse (Ed.), *Advances in man-machine systems research* (Vol. 3). Greenwich, CT: JAI Press.
14. Govindaraj, T., & Vasandani, V. (28 February 1991). *Qualitative simulation and intelligent tutoring aids for training in the operation of complex dynamic systems* (Final Technical Report CHMSR-91-1). Atlanta: Georgia Institute of Technology, Center for Human-Machine Systems Research.
15. Towne, D. M. (1987). The generalized maintenance trainer: Evolution and revolution. In W. B. Rouse (Ed.), *Advances in man-machine systems research* (Vol. 3). Greenwich, CT: JAI Press.
16. *ITS-88: Intelligent Tutoring Systems*. (1-3 June 1988). Montreal: University of Montreal.
17. Sleeman, D., & Brown, J. S. (Eds.). (1982). *Intelligent tutoring systems*. New York: Academic Press.
18. Rouse, W. B. (Ed.). (1987). *Advances in man-machine systems research* (Vol. 3). Greenwich, CT: JAI Press.
19. Psotka, J., Massey, L. D., & Mutter, S. A. (Eds.). (1988). *Intelligent tutoring systems: Lessons learned*. Hillsdale, NJ: Lawrence Erlbaum Associates.
20. Burns, H. L., Parlett, J. W., & Redfield, C. L. (Eds.). (1991). *Intelligent tutoring systems: Evolutions in design*. Hillsdale, NJ: Lawrence Erlbaum Associates.
21. Polson, M. C., & Richardson, J. J. (1988). *Foundations of intelligent tutoring systems*. Hillsdale, NJ: Lawrence Erlbaum Associates.
22. Hamill, B. W. (1990). Three issues for intelligent tutoring. *Machine-Mediated Learning*, **3**, 169-187.

Research Notes

New Method to Reconstruct Hubble Telescope Images

Drs. Michael Cobb and Paul Hertz of The Naval Research Laboratory's E.O. Hulburt Center for Space Research have implemented a system to reconstruct images from the Hubble Space Telescope (HST) using NRL's Connection Machine, a massively parallel computer.

According to Dr. Hertz, it is well-known that problems with the primary mirror of the HST is causing the images obtained from the telescope to be blurred. The blurring prevents all of the light from being focused in one place. As a result, most of the light and most of the information in the image is spread out in a halo around the center of the image. The shape of this "spreading out" is called the point spread function (PSF).

Dr. Hertz reports that certain optical elements of the HST's principal camera distort the PSF and, since the distortion is different for different parts of the total image, a single PSF cannot characterize the distortion for a full HST image; hundreds to thousands of PSFs need to be considered.

Traditional image reconstruction techniques use a single PSF for the entire image and, until now, computer enhancement of HST images has used a single, average PSF for the entire image.

By using parallel processing to reconstruct HST images, says Dr. Hertz, the task can be efficiently divided so that each processor handles its own PSF and reconstructs one small portion of the image. The Laboratory's Connection Machine, which has 16,384 processors and can store 16 billion bits of information in its memory, performs over two billion calculations every second during this reconstruction process. Accordingly, Dr. Hertz reports, an HST image can be reconstructed in under two minutes.

Ground-Based Laser Measures Vibration of Orbiting LACE Satellite

Laser measurements of vibrations in the LACE spacecraft were recently conducted by the Massachusetts Institute of Technology (MIT)/Lincoln Laboratory and the Naval Research Laboratory (NRL) under the sponsorship of the Strategic Defense Initiative Organization (SDIO). The successful tests represent the first time ever that ground based lasers have measured vibrations of an orbiting spacecraft. The measurements were made at the MIT Lincoln Laboratory FIREPOND

Laser Radar facility located near Westford, Mass. The tests involved observations of the vibration induced Doppler signatures of the LACE satellite (object #20496) using the narrowband heterodyne CO₂ laser radar. The tests, conducted during the first two weeks of January 1991, measured tip vibrations in a long deployable/retractable boom during and after retraction.

The LACE spacecraft was built for the SDIO at the Naval Research Laboratory and was launched on February 14, 1990. Its primary mission is to evaluate atmospheric compensation techniques for ground-based lasers. Another primary experiment on LACE is to observe rocket plumes from an on-board ultra-violet tracking camera.

For the purposes of measuring on-orbit system vibrations, the LACE spacecraft has IR germanium retroreflectors mounted on the lead boom, satellite body and the trailing boom. The germanium reflectors serve as targets for the FIREPOND laser radar. The laser radar operated with a peak transmit power of 600 W and a nominal pulse duration of 2 ms (or 3 mm/s velocity resolution at a wavelength of 10.6 microns).

The primary test objective was to obtain vibration data from the lead boom. To this end, the LACE lead boom was retracted from 80 feet to 15 feet. Boom vibration measurements were collected during and after the boom retraction maneuver. The relative motion is due to the fact that the spacecraft is always pointed at the earth, plus boom retraction speed. Preliminary analysis of the laser radar data suggests the presence of vibrations associated with the boom retraction mechanism and with the persistence of lower frequency models associated with system vibration after completion of the boom maneuver. The experiments were repeated and similar results were obtained during satellite passes on January 7, 8, and 10. These measurements demonstrate the ability to obtain useful structural information of space-based platforms from a remote narrowband IR laser radar system. This information is critical to understanding how to design and control large space-based structures.

Mysteries of 'ULF' Frequency Probed by ONR Sound Scientists

Understanding how sound propagates, and how it can be detected in the ocean, is of great benefit in terms of developing future Navy sonar systems.

At the Office of Naval Research, a project is underway to better understand the source generation mechanisms and the space-time characteristics of ambient noise that sonar systems may encounter. The research combines theoretical and numerical modeling with observations along the shore and at sea to verify existing hypotheses and will use the results to develop ambient noise models.

Scientists are attempting to understand fully the generation and propagation of seafloor noise along the east coast and the Arctic in the Ultra Low and Very Low Frequency (ULF/VLF) band (1mHz-50 Hz). The primary focus of the experiment will be in the ULF range.

ULF—meaning Ultra Low Frequency—is a term coined by Dr. Randall Jacobson, one of ONR's geology and geophysics scientific officers, in identifying a frequency range where noise generation and propagation is least known. He chose 1 Hz as the dividing line between ULF and VLF because that frequency band has been hardly researched in the marine environment and because most current hydrophones are capable of reliably making measurements only down to 1 Hz.

"We believe we can make a lot of progress very quickly in the ULF range," says Jacobson, "because of new sensor technology. Most of the ULF-related noise is caused by sea surface phenomena: swell, chop, breaking waves and wave-wave interactions. The noise is so strong that it propagates onto land. Land seismologists long ago mistakenly called this noise 'microseisms,' meaning literally, small earthquakes."

The research is made up of four primary field programs: SAMSON (Sources of Ambient MicroSeismic Oceanic Noise); ECONOMEX (Environmentally controlled Oceanfloor Noise Monitoring Experiment); and BASIC (Beaufort Ambient SeismoAcoustics under Ice Cover) and NOBS (Noise on Basalts and Sediments)..

The objective of BASIC is to determine the contribution of long-range microseismic noise to that locally generated. It is also designed to search for ULF noise that is specific to the presence of shore fast ice.

NOBS' objective is to better understand the impact of bottom type (e.g. thickly sedimented vs. hard rock) upon noise propagation and scattering.

The objective of SAMSON is to understand the source generating mechanisms of ULF noise at the coast line and to determine propagation characteristics of the noise into the continent and out into the ocean.

The objective of ECONOMEX is to understand how ULF and VLF noise is generated along the continental shelf-slope break, its distribution with depth, and its propagation characteristics under wider environmental conditions than SAMSON.

SAMSON is designed to examine the fine-scale details of ULF ambient noise during an intensive, six-week period in the fall, whereas ECONOMEX is designed to examine the synoptic, seasonal variation of both ULF and VLF noise during a six month winter period. ECONOMEX should provide noise data that overlaps and complements SAMSON. Both programs will be

conducted in conjunction with SWADE (Surface Waves Dynamics Experiment) Program, the best sea surface observational program yet undertaken, and also funded by ONR.

The goal of this Navy research initiative on ambient noise will be realized through the collection and detailed analysis of continental, seafloor and mid-water seismo-acoustic data and surface environmental data from shallow to deep water in a transect extending from the North Carolina coast to ocean basin depths.

A linear network of seafloor seismo-acoustic instruments recorded the evolution of the seafloor noise field as a function of distance from the coastline and depth of water, while several small coherent arrays were used to study the various components of ambient noise in the ULF band.

"This experiment is deceptively simple in concept," says SAMSON chief scientist John A. Orcutt from the Institute of Geophysics and Planetary Physics, Scripps Institution of Oceanography, University of California, San Diego, The experiment called for a line of seismo-acoustic instruments beginning at the coast of North Carolina and extending well offshore beyond the shelf break.

"We augmented this array with several coherent arrays to study the various components of ambient noise in the ULF bank," Orcutt explains.

This shallow water to deep water experiment addresses several scientific problems which cannot be answered with an experiment which is sited only in deep water. The transition from shallow to deep water requires very large changes in the modes of propagation of seismo-acoustic waves. Consequently, the scientists believe the study of this transition will teach them much more about the partitioning of energy in oceanic ambient noise between modes than the collection of measurements solely at deep or shallow water sites.

"The relative importance of different sources of ambient noise must also change dramatically from shallow water to deep water," says Orcutt, "if only because the shelf has a strong effect on the propagation of surface gravity waves. We will address the problem of the generation of primary frequency seismic waves (microseisms) which are known to be generated near the coast, and the excitation of seafloor noise by both nonlinear (deep water) and linear (shallow water) processes will be contrasted," he said.

A component of SAMSON used about 37 land-based seismometers in an array in North Carolina. Because of the massive amount of data collected, it is essential that the data be edited and some preliminary analyses be performed in the field. For that purpose, he hopes to have a high-level field computer available which will permit the scientists to employ a fairly sophisticated level of analysis in the field. "An important advantage of having extensive analysis capability at the experiment site," says Orcutt, "is that it allows us to change recording parameters or instrument sites if suggested by the data."

An array of bottom-mounted pressure sensors placed offshore near Duck, North Carolina, and operated by the Army Corps of Engineers, provided a continuous measure of the

surface gravity wave field propagating over the very inner shelf. This array measured the directional spectrum of wind-generated surface gravity waves and infragravity waves. "The data (gathered) hopefully will allow us to understand the responsible mechanisms," says Orcutt. Because of the large amount of quality data collected, it will take a couple of years of analysis to fully understand and develop a predictive model of noise generation and propagation.

When the Office of Naval Research initiated the research, it realized its source of instrumentation was limited. Few, if any, instruments of the type needed to make appropriate measurements existed.

According to Jacobson, ONR felt not enough instruments were available "to make array-kinds of measurements that would allow us to tell what direction the energy was propagating or how coherent it might be from sensor to sensor." So the Marine Geology and Geophysics Program funded approximately \$1 million in fiscal year 1988 and almost \$2 million in FY89 to build a suite of new instruments. These included 31 state-of-the-art ocean bottom seismometers, about 10 differential pressure gauge single sensor instruments, as well as a small collection of related instruments.

"Those instruments," says Jacobson, "were built by WHOI (Woods Hole Oceanographic Institution), Scripps Institution of Oceanography, the University of Washington and Massachusetts Institute of Technology. They are now coming on line for use and are part of the seismic experiment."

As this ONR ULF/VLF initiative gets underway on site, several payoffs are anticipated. Among the most promising payoffs include developing means to predict noise fields and to develop models, optimal sensors and configurations for sonar surveillance arrays. ONR also hopes to improve signal detectability in poor signal to noise areas, which will be of great benefit to future sonar development.

Sound Localization In Human Audition:

Auditory psychophysicists long have recognized that localization of sound sources along the horizontal dimension (azimuth) is computed in the auditory system on the basis of differences between the ears in sound intensity and arrival time. However the mechanism by which the elevation of sources is computed is less well understood. One hypothesis is that elevation is determined somehow by the spectral filtering characteristics of the pinna (outer ear). Strong support for this hypothesis has now been obtained in experimental and computational work carried out at the University of Florida. The initial finding was that while narrow and broad-band sounds are localized in elevation much more accurately than narrow. Later experiments revealed systematic center-frequency dependent errors in elevation judgments by human listeners. For example: (1) when presented at a -20° elevation, a narrow-band sound centered at 6kHz will be judged as located at an elevation of +50°, and (2) a 12 kHz band presented at +30° will be judged as coming from 0° elevation. In experiments that followed, a small microphone was

inserted into the external ear canal (meatus), and flat spectrum broad-band sounds were presented from varying elevations. The important finding was that the measured intensity of each spectral component of the incident sound recorded at the meatus was related systematically to the elevation of its source. Thus, for example, components in the 6kHz region show a higher recorded intensity at +50° elevation relative to other frequency bands, and components in the 12kHz region show a relatively higher intensity at 0°. These findings led to a formal model in which elevation is computed from the empirically derived transfer function of the individual human pinna: in effect, elevation is computed by the model directly from the spectrum of the sound reaching the middle ear. The model accurately predicts the judged location of both narrow and broad-band sounds, and fully accounts for individual differences in localization accuracy.

Gamma-Ray Instrument Launched

Scientists at the Naval Research Laboratory (NRL) have developed the Oriented Scintillation Spectrometer Experiment (OSSE), one of four large instruments on the National Aeronautics and Space Administration's (NASA's) Gamma Ray Observatory (GRO), which was launched on April 5. The GRO is designed to operate for a minimum of two years.

GRO was carried into space by the space shuttle Atlantis and is the second of four "Great Observatories" scheduled for launch by NASA. The first, the Hubble Space Telescope, was launched in 1990. The Advanced X-ray Astrophysics Facility (AXAF) and the Space Infrared Telescope Facility (SIRTF) are scheduled to be launched in the late 1990s. NRL investigators note that these four major missions will cover nearly the entire electromagnetic spectrum, from infrared to high-energy gamma rays, providing a major advance in the ability of scientists to understand the universe.

OSSE is designed to measure the spectra and time variability of gamma rays from sources as near as the sun to active galaxies a billion or more light years away. OSSE is expected to detect gamma ray line emissions and the continuum spectrum of celestial sources with 10 to 20 times better sensitivity than previous instruments.

The scientific objectives of OSSE include studies of exploding stellar objects, such as

- supernovae and novae (the sites where heavy metals are believed to be created);
- the study of very compact objects like neutron stars, black holes that may exist in space, and pulsars;
- investigation of powerful objects in the centers of distant galaxies;
- a survey of the galactic plane and the interstellar medium;
- observations of solar-flare gamma-ray and neutron emissions; and
- a partial sky survey.

OSSE uses four gamma-ray detectors made from materials that "scintillate," i.e. emit visible light, when gamma rays interact with the materials. Light detectors can then be used to count the gamma rays and measure the energy of each incoming gamma ray in the energy range from 0.1 to 10 million electron volts. This energy range includes the region where most nuclear processes occur and which are characterized by the emission of nuclear gamma rays. Gamma ray spectra are characteristic of the specific nuclear species from which they originate.

Measuring the spectra of these gamma rays will enable NRL scientists to determine the nature of the nuclear processes occurring in these very energetic celestial sources in much the same way as optical spectra enable scientists to study phenomena in ordinary materials and the surfaces of stars. This capability will provide new insights into the physical processes occurring in exploding celestial objects, such as the supernova, which occurred in 1987 in a nearby galaxy (the Large Magellanic Cloud), collapsed objects such as suspected black holes and neutron stars, and the very energetic processes that power distant active galaxies, which emit millions of times the energy of our own Milky Way. NRL scientists note that the instrument will also provide measurements of the gamma-ray environment in low-earth orbit where future scientific missions will operate.

Overall scientific and technical development of OSSE has been under the direction of Dr. James D. Kurfess, principal investigator for the experiment and head of the Gamma and Cosmic Ray Astrophysics Branch of NRL's Space Science Division. Other members of the team include other scientists from NRL, Northwestern University, Clemson University, and the Royal Aerospace Establishment in England. The design and development of OSSE were completed under contract with the Ball Aerospace Systems Division in Boulder, CO.

Following completion of the development of OSSE, the instrument was delivered to TRW, Inc. in Redondo Beach, California. For 15 months, OSSE and the other three experiments underwent integration onto the GRO spacecraft and extensive testing in preparation for its launch.

Manipulation of Atoms:

Joseph Strosio and coworkers at the National Institute for Standards and Technology supported in part by ONR, have recently demonstrated the ability to manipulate individual atoms adsorbed on surfaces using a scanning tunneling microscope (STM). The experiments were performed on clean GaAs and InSb surfaces obtained by cleaving in ultra-high vacuum. Submonolayer coverages of Cs were deposited on these substrates held at room temperature. As in their previous work, it was shown that the Cs atoms tend to form long zigzag chains which can be imaged by the STM operating at negative tip biases. In this recent work it was demonstrated that a Cs atom could be induced to diffuse on the surface by applying a short positive voltage pulse, typically 1 to 3 V for 0.3 to 1 s. By this

technique one can create regions of nearly complete coverage starting with an average coverage of less than 10 percent. One can also create clusters of atoms which may have sufficient density to exhibit metallic properties. This technique appears to be fairly general and should work with a variety of metal atoms deposited on semi-conductor and metal surfaces, opening up the possibility of creating novel nanostructures not found in nature. Such nano-structures may have unusual electronic, magnetic, and chemical properties which could be exploited in future applications.

New Process for Growing Diamond Crystals

A team of scientists at the Naval Research Laboratory's (NRL's) Optical Sciences Division reports the discovery of a new high-temperature epitaxy process for growing millimeter-sized, transparent diamond crystals using an oxygen-acetylene flame.

Dr. Keith A. Snail and Dr. Leonard M. Hanssen, members of NRL's Accelerated Research Initiative (ARI) on diamond, say the process represents a major breakthrough in the field of low-pressure diamond synthesis.

The ability to grow large high-quality diamond crystals at temperatures above 1250°C was discovered at NRL over 2-1/2 years ago. At that time, substrate temperatures above 1100°C were normally thought to favor graphite, rather than diamond growth, but the NRL scientists now report that this conception is not completely accurate.

The key elements of the discovery involve heating a natural diamond seed crystal to temperatures of 1200-1500°C in an oxygen-acetylene flame produced by a commercial welding torch. During one-hour depositions, a faceted single diamond crystal can be grown on top of the seed crystal at a rate of over 150 microns/hour. The acetylene provides the carbon source for growing the diamond crystal, as well as a high hydrogen flux for keeping the diamond's surface from turning to graphite. The cost of the capital equipment and gases is extremely low and scaling to larger areas appears achievable.

In a related discovery, Dr. Snail, Dr. Cheinan Marks, and a team from the University of Minnesota, led by Professor Emil Pfender, report that a DC plasma torch can also be used to grow millimeter-sized diamond crystals at high temperatures (1200-1400°C). The growth rates observed (200 microns/hour) are the highest ever reported for the epitaxial synthesis of macroscopic diamond crystals at low pressures. Dr. Snail anticipates that, with continued research and development, NRL's discovery will eventually enable the growth of carat-sized crystals and boules of diamond.

The long-term benefits of this discovery could be applied to several areas. First, since diamond's thermal conductivity is five times that of copper at room temperature, large diamond crystals would be useful as heat sinks for ultralarge scale integrated (ULSI) electronic circuits, high-power electronic devices, and certain laser diode arrays. In each of these appli-

cations, diamond's excellent ability to disperse heat would be used to greatly reduce temperature fluctuations that can adversely affect the performance of these devices. Second, Dr. Snail notes that diamond is a semiconductor with many unique advantages over silicon. NRL's high-temperature epitaxy process may spawn a range of diamond electronic devices and technologies, including high-temperature, high-power switches; blue light emitting diodes (LEDs); diamond boules and wafers; radiation-hardened electronics, as well as unforeseen devices.

Currently, the diamond grown with NRL's high-temperature epitaxy process is not gem quality; however, by varying the growth conditions, nearly defect-free diamonds can be grown, but at reduced growth rates. The NRL team is scheduled to present details of this process at an international diamond conference in France this coming fall.

Persian Gulf Circulation Model

Dr. Lakshmi Kantha while employed at the Naval Oceanographic and Atmospheric Laboratory assisted the Naval Oceanographic Office in adapting a primitive equation ocean circulation model to the Persian Gulf. The model is a version of the Mellor shallow water model which has performed reasonably well in Delaware Bay, New York Harbor and the Mid-Atlantic Bight. It has recently been used in data assimilation experiments in the Gulf of Mexico and the Atlantic Bight. The Persian Gulf version is the first model to provide an operational fleet support product on the Navy CRAY YMP large scale computer located at Bay St. Louis, Mississippi. The model is being forced by regional wind fields and fluxes generated by the Fleet Numerical Oceanographic Center, Monterey, California and astronomic tidal inputs.

Informal feedback to the Commander, Naval Oceanography Command staff has indicated that the model was useful in assessing mine drift and oil slick movement. Model improvements are currently in progress. A quantitative assessment of model performance has been initiated by the Naval Oceanographic Office, Bay St. Louis, Mississippi. Dr. T. Kao of Catholic University and Dr. S. Chao of the University of Maryland recently briefed the Office of Naval Research on

their Persian Gulf modeling research. They found that buoyancy fluxes, resulting from evaporation, establish and maintain a thermohaline circulation that can be as important as the wind forced circulation. The Gulf, because of its shallowness, is strongly influenced by wind forcing. The interannual effects of the wind on circulation are being examined. In addition, the response of the Gulf to tidal and wind set up forces are being studied. The assimilation of available data such as satellite IR may be helpful in predicting general circulation and is being explored. This is the first of a series of enclosed sea models that are being developed and/or improved for operational use.

First Demonstration of Atom Interferometry

Professor D. E. Pritchard at the Massachusetts Institute of Technology has reported the first demonstration of atomic interference in an atom interferometer. A 3 grating geometry was used in which the interfering beams were separated in both position and momentum. A highly collimated beam of sodium atoms having a de Broglie wavelength of 0.16 Angstroms and high-quality 0.4 micron period gratings were used. In these first experiments, the interferometer's phase was determined to a precision of 0.1 radian using one minute averaging time. This work was supported by the Office of Naval Research. Atom interferometry has potentially important applications to improved gyros, geodesy, gravity gradiometers, tunnel detection, and as an ultra-sensitive diagnostic probe for fundamental physics measurements. The real significance of these results is not that atom interferometry is scientifically possible. This was expected based on the previous demonstration of interferometry using non-zero rest mass particles (as compared to photons) such as neutrons and electrons. In order to observe stable atomic interference fringes, various requirements on mechanical stability and alignment must be satisfied. Specifically, these requirements are due to the sensitivity of the interferometer to rotation, acceleration, and translation of the gratings. This work demonstrated for the first time that these requirements can be achieved in practice.

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

OFFICIAL BUSINESS
PENALTY FOR PRIVATE USE, \$300