PATRICK T. NEWELL, SIMON WING, CHING-I. MENG, and VINCENT SIGILLITO

A NEURAL-NETWORK-BASED SYSTEM FOR MONITORING THE AURORA

Electrons and ions from the various regions surrounding the Earth rain down upon, or precipitate into, the upper atmosphere at high latitudes. The visible manifestation of this precipitation is the aurora, which in the Northern Hemisphere is also called the northern lights. The particles precipitating into the Earth's atmosphere can be used as a diagnostic—a "space weather" monitor. A series of Air Force Defense Meteorological Satellite Program satellites, collectively in continuous operation, measure precipitating particles in the range where most of the energy flux is carried. We have developed an algorithm for automatically monitoring the high-latitude precipitation, including a neural-network–based identification of the source region of all precipitation observed. The result is an enormous and sophisticated database, one use of which is to determine the appropriate mapping of ionospheric magnetic field lines into near-Earth space. The database is also being provided on a limited basis to the space physics community as a service. In addition, the system is a logical step in the development of a real-time capability to predict space weather.

INTRODUCTION

A continuous outflow of hot ionized gas from the Sun—the solar wind—blows against the Earth. The Earth's magnetic field forms a bubble that keeps most of the solar wind out. Within this bubble, called the magnetosphere, various regions of trapped or quasi-trapped plasma exist. The charged particles that make up the plasma, mostly electrons and protons, are guided by the Earth's magnetic field lines, and a portion of them are continually striking the Earth's upper atmosphere at high latitudes in a ring around the magnetic poles that is called the auroral oval. Figure 1 shows an ultraviolet image of the polar region from an altitude of about 1000 km; approximately three-quarters of the auroral oval is clearly visible. The image is from the Air Force Polar BEAR (Polar Beacon Experiment and Auroral Research) satellite, built at APL.

The magnetic field lines throughout the magnetosphere are focused together above the Earth's polar regions, since the field lines of a dipole converge into the poles. Hence, a low-altitude polar-orbiting satellite, when suitably equipped with the appropriate particle detectors, can rapidly sample the various plasma regimes surrounding the Earth. This sampling ability makes measurements of the auroral oval a logical candidate for monitoring "space weather," that is, the state of the magnetosphere. In effect, the polar regions provide a projection-screen view of near-Earth space.

The upper atmosphere of the Earth acts as the screen; the precipitating electrons excite the airglow display known as the northern and southern lights (aurora borealis and aurora australis, respectively). Although many scientific satellites have investigated the auroral regions at various times, by far the largest and most comprehensive data set of auroral particle measurements comes from the Air Force Defense Meteorological Satellite Program (DMSP). The DMSP satellites are operated continuously, providing a completeness of record unavailable elsewhere. The DMSP data set is thus the best choice for routinely monitoring the auroral oval and other high-latitude precipitation, and is therefore a fairly good indicator of the state of the magnetosphere.

Although our primary interest is in basic research, we have found some practical use in monitoring space weather. At high latitudes, intense auroral activity can knock out power stations and disrupt radio communications and over-the-horizon radars (looking over the polar cap), and correlated processes at high altitude can charge spacecraft to high voltages. Moreover, in looking toward a future of increased presence in space, it makes sense to understand the Earth's space environment as much as possible.

Therefore, our research group has been pursuing a campaign in which we (1) try to identify as accurately as possible the magnetospheric (or solar) source of all forms of particle precipitation observed at high latitudes, and (2) teach a computer to perform the same identifications. We have, in effect, been developing a system to monitor space weather routinely with the following goals in mind: (1) providing a service to the space physics community (providing other magnetospheric researchers with the position and state of the auroral oval); (2) performing statistical manipulations on the data set, leading, for example, to a more accurate mapping of the ionosphere to the magnetosphere; and (3) performing routine monitoring, which is an important step in developing a predictive capability.

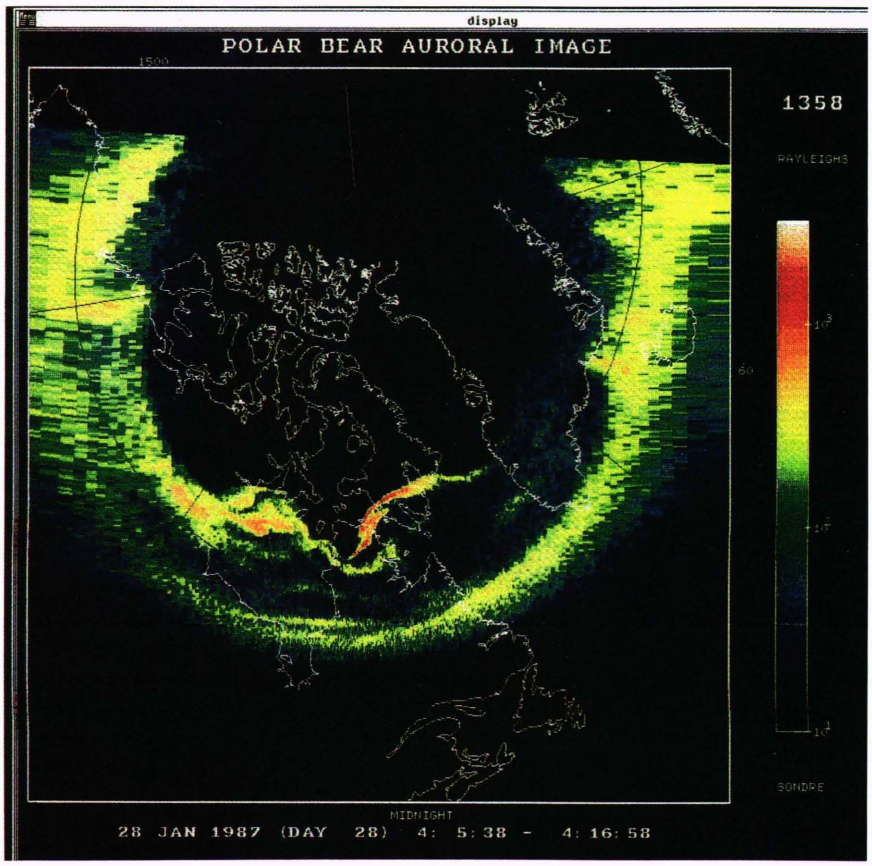

Figure 1. An ultraviolet image of the Earth from the Polar BEAR satellite at an altitude of about 1000 km. The false color scale gives the intensity in Rayleighs of auroral oval emissions at 135.8 nm. About three-fourths of the auroral oval is clearly distinguishable. The thickening at the sides is due to the slanted viewing angle when looking away from nadir. For orientation, the red line at the top points toward the Sun.

PRECIPITATING PARTICLES AND THE DMSP DATA SET

The DMSP satellites are all polar-orbiting satellites in nearly circular polar orbits at an altitude of 835 km and an orbital inclination of 98.7°. Normally, two DMSP satellites are in operation at any given time. Currently, F8 and F9 are operational; the previous pair was F6 and F7. The DMSP satellites are sun-synchronous; F6 and F8 are in approximately the dawn/dusk local-time meridian, and F7 and F9 are in the prenoon/premidnight local-time meridian. The instruments we use are the SSJ/4 electrostatic analyzers, which measure electrons and ions once per second from 32 eV to 30 keV in twenty logarithmically spaced steps.[1] The DMSP satellites are three-axis stabilized, and the particle detector apertures are always directed toward local zenith. At the magnetic latitudes (MLAT) of interest to auroral researchers (say, greater than 50° MLAT), this condition means that only particles aligned approximately in the direction of the Earth's magnetic field are observed. Such field-aligned particles will hit the Earth's upper atmosphere and deposit their energy therein. Particles less field-aligned will be reflected from the magnetic mirror of the Earth's converging field lines and bounce back into space.

We classify precipitating particles in the energy range measured by DMSP as having come from one of seven sources. Because charged particles tend to follow magnetic field lines, these various source regions can be mapped from distant locations to the polar region ionosphere along converging magnetic field lines to DMSP altitudes (835 km) (see Fig. 2). The source regions are, in order of their latitude of most likely occurrence, from lowest to highest: photoelectrons, created by the more energetic portion of the solar spectrum; the central plasma sheet (CPS), high-energy precipitation so called because it is conventionally believed to map to the Earth's central plasma sheet or its dayside extension;[2] the boundary plasma sheet (BPS), conventionally believed to map to the boundary between the Earth's plasma sheet and the lobes; the low-latitude boundary layer (LLBL), a region just inside the Earth's magnetosphere along the flanks; the cusp, a limited region at high latitudes on the dayside where the shocked solar wind has fairly direct access to low altitudes above the Earth; the mantle, a region of de-energized solar wind, above and below the magnetosphere; and polar rain, soft electrons also of solar origin entering from deeper down the tail.[3]

Of these seven regions, three—photoelectrons, the CPS, and polar rain—are easy to identify because they have the least spectral and spatial fine structure; that is, these regions tend to be comparatively homogeneous. The cusp was first identified by Heikkila and Winningham[4] in 1971; more recently, identification of the cusp has been made quantitative and systematic,[5–7] and this region can now also be considered to be reliably identifiable in a data set of precipitating particles. More work needs to be done in establishing firmly that the LLBL and the BPS can be distinguished reliably on the basis of only the spectral properties of the precipitation; but the lower

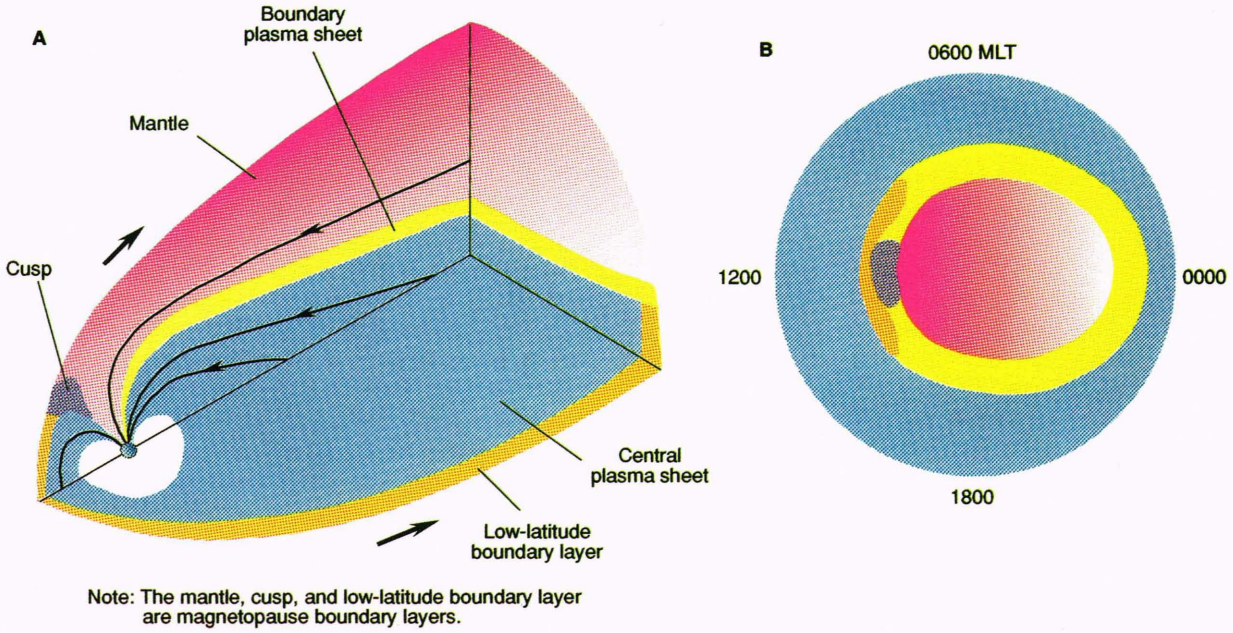

Figure 2. Mapping the magnetosphere into the ionosphere along magnetic field lines. **A.** Schematic diagram of various observed magnetospheric boundary layers. **B.** Schematic diagram of the regions occupied by the various boundary layers mapped into the ionosphere. (Reprinted, with permission, from Vasyliunas, V. M., "Interaction between the Magnetospheric Boundary Layers and the Ionosphere," in *Proc. Magnetospheric Boundary Layers Conf.*, European Space Agency Special Publication, ESA SP-148, 1979. Published by courtesy of the European Space Agency.)

energy of the electrons and especially the ions in the LLBL would seem to make such identification practical. The plasma mantle has only recently been identified at low altitude,[8] except for the obvious case of directly poleward of the cusp. It is outside the scope of this article to review each of these identifications in detail, although we will give some examples of the identifications later.

AUTOMATED IDENTIFICATION OF HIGH-LATITUDE PRECIPITATION

The system for automated identification of high-latitude precipitation consists of a feed-forward, backward-error-propagation neural network (as explained later) and a post-processor. The neural network operates on each second's worth of data, namely, one electron and one ion spectrum. Thus, the network is making decisions without context, which is the most serious limitation. The output of the neural network goes to a simple expert system—a set of explicit rules for converting the stream of up to 1500 one-second identifications into a comparatively small number (<10) of distinct regions encountered. We consider each of these steps in the sections that follow.

The Neural Network

Interest in neural networks has recently surged; conferences are being held and journals devoted to the subject are appearing. The underlying ideas are not recent, however; theoretical forerunners reach back at least to the 1940s.[9,10] Very simple networks known as perceptrons[11] were popular in the artificial intelligence community in the late 1950s and early 1960s, until their severe limitations became clear.[12] Although it was realized at that time that more complicated networks could be constructed, no learning algorithm for these networks existed, so that interest waned.

More recently, it has become clear (e.g., Ref. 13 and references therein) that an efficient method, known as the generalized delta rule and based on the method of steepest descent, could be used to "train" a network to learn the values for weights between nodes. Spectacular successes with this method, such as a network that taught itself to pronounce English,[14] have led to the resurgent interest.

Recently, interest has moved beyond the artificial intelligence community, and neural networks have been applied to scientific and technical problems, particularly in the fields of signal and image processing.[15,16] At APL, Sigillito et al.[17] have implemented a neural network to identify good and bad radar returns in ionospheric studies, although the method has never been adopted for practical use. Other examples of neural network applications in image and signal processing can be found in the July 1988 issue of the *IEEE Transactions on Acoustics, Speech, and Signal Processing.*

As the name implies, the idea of a neural network is to simulate the distributed processing of the human brain. Several models of neural networks are under study; the one we have chosen, which is the most commonly used, is a feed-forward, backward-error-propagation network. It contains several layers of nodes; each layer accepts input only from lower layers of nodes and

outputs only to the next higher layer of nodes. For definiteness, we will now describe the particular network we used in our application.

Our network had three layers: the input layer, the output layer, and one intermediary or hidden layer. We tried networks with two hidden layers, but with little or no improvement in accuracy. The input layer was composed of forty nodes, consisting of the thirty-eight particle count rates (a nineteen-point electron and nineteen-point ion spectrum), Magnetic Local Time (MLT), and MLAT. Eight output nodes were included, one for each type of dayside precipitation that we wanted identified: photoelectrons, the cusp, the near cusp, LLBL, BPS, the dayside extension of the CPS (also known as the outer radiation zone), the mantle, and polar rain. The output nodes produced "1" if the type of precipitation they specialized in was present and "0" otherwise.

Each of these forty input nodes is connected to each of the hidden nodes in the second layer (a hidden node is any node that is not an input or an output node). For the dayside network, thirty hidden nodes gave the best results; for the nightside network, only fifteen were necessary. These, in turn, were connected to each of the eight output nodes. The exact number of hidden nodes to use is a matter of judgment and experimentation; we tried various values before settling on the final numbers.

Signals propagate from layer i to layer $i + 1$ within the network as follows: The output O_j^{i+1} of node j of layer $i + 1$ is related to the outputs of layer i by the rule $O_j^{i+1} = 1/[1 + \exp(-\Sigma W_{ij} O_j^i)]$, where the summation is over all nodes j within a layer. The network is defined when the weights W_{ij} are determined for each node j within each layer i.

The technique for determining these weights is to train the network by using a set of example inputs and target outputs. Initially, the weights are selected at random, and the outputs will differ greatly from the target answers. The rule for adjusting the weights, known as the generalized delta rule, is based on the method of steepest descent.[13] Therefore, the change in the weights ΔW_{ij}, after each target output is compared with the actual output, is given by $\Delta W_{ij} = \eta \partial E/\partial W_{ij}$, where the error function $E = \Sigma_j (t_j - O_j)^2$ is based on the difference between the target outputs t_j and the actual outputs O_j. The constant η is called the learning rate and is typically chosen to be in the range from 0.1 to 0.9. For our large data set, an η value of 0.08 or less proved necessary. In practice, it is often advantageous to add a "momentum" term that is proportional to the previous ΔW_{ij} to help prevent oscillation about local minima (our momentum coefficient was 0.3). The corrections to the weights are propagated backward, starting from the output layer and working toward the input layer. Hence, the network feeds signals forward from the input layer to produce its output and propagates error correction backward on the basis of closeness of the actual outputs to the target outputs.

As a practical matter, therefore, the work required by the scientist in training the network to identify features of interest is simply to come up with a set of example inputs and target identifications. Typically, 50 to 100 examples of each signature to be identified are required, normally including roughly the same number of examples having the phenomena present as having the phenomena absent. These examples and targets are presented repeatedly until the network converges to identifying a high percentage of the training set.

It is then necessary to test the network on data it was not trained on, to ensure that the network has truly learned to generalize and has not simply memorized the correct answers for the limited training set. This second data set is called the testing set. In our case, we used training sets for the dayside and nightside of sizes 100 and 63, respectively, and a testing set of size 10 for both networks. A larger training set was necessary for the dayside than for the nightside simply because more classification categories were used for the dayside. The nightside network reached 98% correct on the training set and 96% correct on the testing set. The dayside network, with its larger number of categories and finer distinctions, was 91% correct on the training set and 90% correct on the testing set. Recall that these are one-second identifications, so that most of the incorrect identifications are due to fluctuations within a region. Most mis-identifications can be corrected by post-processing, as described in the next section.

Rules-Based Post-Processing

Regions of precipitation extend from a few consecutive seconds to several hundred seconds; an entire polar pass can be up to 1500 seconds, where each second includes a thirty-eight point spectrum and two location coordinates. Rather than attempting a neural network with 60,000 input nodes, the computation involved in the network was kept to a manageable level by having the network evaluate a single spectrum (one second of data) at a time. Thus, a second level of processing is needed to convert the stream of one-second (one-spectrum) data into regions of precipitation, each of which has one overall identification. Most mis-identifications can be corrected at this step on the basis of contextual information. A set of explicit rules, in effect, a simple expert system, was chosen for the second-stage processing. Some of the necessary rules are obvious; for example, a string of identifications such as BP, BP, BP, CP, BP, BP, BP is stored as one region of BP, signifying BPS. Others involve knowing something about the magnetospheric configuration, for example, that CPS cannot be found poleward of the cusp, although unusually energetic polar rain might be. Altogether, about thirty such rules were adopted, each of which resolves a conflict or unlikely string of identifications according to various quantitative conditions.

Although it would be prohibitive to describe each of the rules in detail since some are quite complicated, the flavor of the rules can be shown with a few simple examples. Three or more consecutive seconds that have the same identification, say x, are considered to be a region R_x. One or two seconds of an identification is taken to be a short region S_x. Thus, the general version of the rule first mentioned is $R_x S_y S_z R_x \rightarrow R_x$, that is, the two short regions are absorbed. Another example is that

a string that alternates between regions of polar rain and void, each region of length less than ten seconds, is listed as a single region of either polar rain or void, according to which has the greater number of seconds identified (regions of longer consecutive similar identifications are not combined).

It was found useful to construct a "guess function" to help resolve conflicts in identifications. The guess function is an explicit algorithm for making a reasonable (context-free) guess based on quantitative information about a region (especially integral parameters, such as average energy and energy or number flux). In general, the guess function does not perform as well as the neural network, but it can act as a valuable second opinion in resolving conflicts. For example, a region that satisfies the Newell and Meng[5] quantitative criteria for the cusp proper has a guess value of CU. Consider a case in which the neural network identified a region of LLBL at a higher latitude than the cusp. This identification conflicts with likely magnetospheric morphology, unless the satellite is changing significantly in local time as well as latitude. The resolution would depend on the guess values of the two regions and the lengths of the two regions (regions larger than nine seconds are almost never changed, whatever the conflict). Experience shows that the net effect of the rules is to greatly improve the overall accuracy of the identifications.

We have not made a systematic effort to quantify the overall accuracy of the entire system. We have, however, compared the network performance to our own identifications for more than 100 polar passes. Our experience is that the automated identifications are almost always basically correct for clearly unambiguous regions, although the exact position of the boundary between regions may differ slightly from the one that a human using context would choose. For the most ambiguous situations, the automated identifications are at least tenable. The identifications are least reliable, however, when the precipitation patterns are the most unusual. Unfortunately, the more unusual the magnetospheric conditions, the more likely the period is to be of interest to space researchers. Even under these circumstances, the database of automated identifications can be of some use to other members of the space physics community, for example, by identifying the general position of the auroral oval, even though the detailed region identifications will be less reliable.

The On-Line Database

The entire set of DMSP F7 particle data, as well as data from F9 through the present (with about a six-month lag), has been processed through the neural-network–based system. After the identifications of a given pass are finalized, they are added to the compressed database. The region identification, ephemeris information, and spectral information for the region are stored. The points of peak electron and ion energy flux over the period of the region, based on a sliding three-second average, if the region is longer than three seconds, are also recorded. This information is all stored onto a write-once read-many (WORM) compact disk; the entire database easily fits onto one WORM.

A permanent on-line database is also created; it differs from the main database only in that the detailed spectral information (the count rates in each of the thirty-eight channels) is dropped, thus reducing the storage requirements by about 60%. Integral information such as average energy and total energy flux is retained. In the reduced form, each region identified takes 56 bytes (56 B). The data are stored in one-month files, which average about 1200 blocks (600 kB). For our own use, and for the use of other magnetospheric researchers at APL, instant access to data from any time over the past several years is possible. Moreover, a simple database search algorithm allows one to search by geographic, magnetic,[18] or temporal coordinates, or certain other conditions. This ability should be particularly useful for collaborating with ground-based researchers since it allows a determination of coincidences (overflights of the satellite through the ground-based observer's field of view).

Space physics researchers outside APL are provided with a more restricted access to the data set through the Space Physics Analysis Network (SPAN). Even though no classified research is done on the APL computers connected to this network, it was deemed prudent not to allow outside users to actually initiate processes on the APL Vax computers devoted to space physics. Moreover, to allow the authors exclusive access to statistical results from the network, it is desirable to restrict outside users to forty-eight hours' worth of data at a time. Thus, a computer-based mail system for making requests was developed. The processing of the requests is automated; a program starts up three times per day and responds to the data requests. The procedure for making a request is simple. A rigidly formatted mail message, with the subject being "DATA_REQUEST," is sent to the SPAN address "APLSP::OVAL." The message is three lines long:

starting;19xx,mm,dd,hh
ending;19xx;mm,dd,hh
{return SPAN address}

Space physics researchers interested in the database can find further information published elsewhere.[19]

SOME REPRESENTATIVE EXAMPLES OF NEURAL-NETWORK-BASED IDENTIFICATIONS

Two examples each from DMSP F7 nightside and dayside crossings of the auroral oval are shown in Figure 3. The spectrograms are a convenient method of three-dimensional plotting. The *x*-axis is the spacecraft trajectory and is thus labeled with time and positional data, and the *y*-axis is the energy of the particles being plotted. The color scale gives the differential energy flux of electrons and ions precipitating into the atmosphere. The orange labels under the spectrograms are generated automatically by the neural-network–based system described in the previous section.

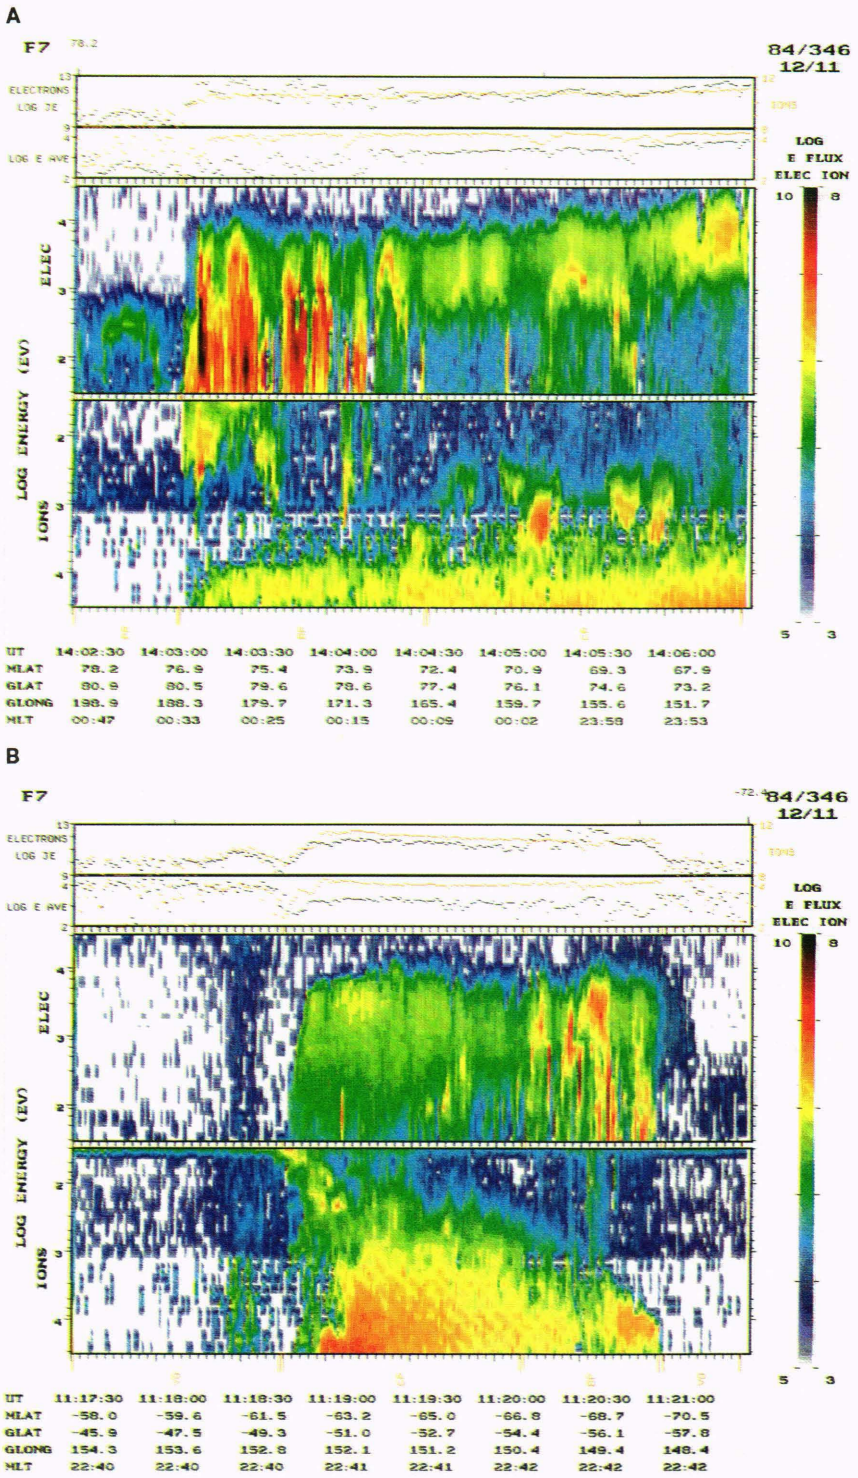

Figure 3. Spectrograms from DMSP F7 showing differential energy flux (eV/cm^2·second·steradian·eV) from 32 eV to 30 keV. The ion energy scale (eV) is inverted (all numerical labels are log values). The orange labels under the spectrograms are generated automatically by the neural-network–based system. The meaning of each label is given in the text. **A.** A nightside crossing with easily distinguishable diffuse and discrete regions of auroral precipitation. **B.** A nightside crossing where the diffuse and discrete regions are less sharply distinct. The counts at 1118:30 UT are due to penetrating radiation belt electrons and are correctly discounted by the network. **C.** (next page) A dayside oval crossing showing a complicated structure, including the cusp proper, low-latitude boundary layer, boundary plasma sheet, and central plasma sheet. **D.** (next page) A dayside oval crossing with a much simpler structure; only the diffuse and discrete auroral regions are seen.

Figure 3A shows a nightside crossing from 11 December 1984 at about 1400 UT during auroral substorm activity. In this crossing, three regions are clearly distinguishable (from right to left): the diffuse, comparatively unstructured aurora (CP); the region of discrete aurora (BP); and polar rain (PR). We believe that few magnetospheric researchers would greatly disagree with the diffuse/discrete auroral boundary selected by the network.

A more challenging example of a nightside crossing through a less active auroral oval is shown in Figure 3B, also from 11 December 1984, at 1117 UT. Again, the identifications are quite satisfactory. The neural network successfully discounted spurious counts due to penetrating electrons from the outer radiation belt at 1118:30 UT; this region is correctly labeled as being void (VO) of precipitation in the 32 eV to 30 keV range. Although instances occur when it is difficult for anyone (or any-

C

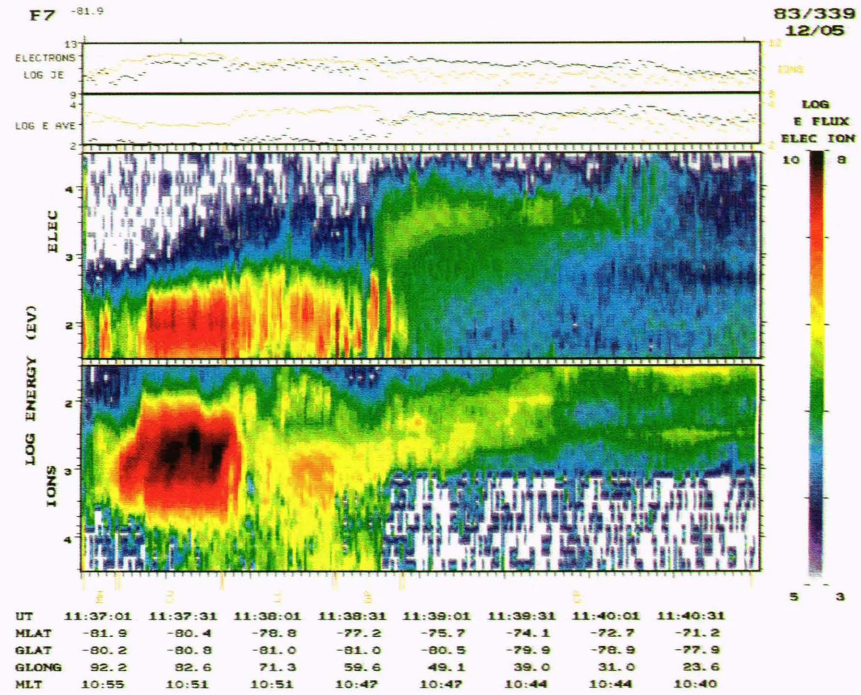

Figure 3. (continued).

D

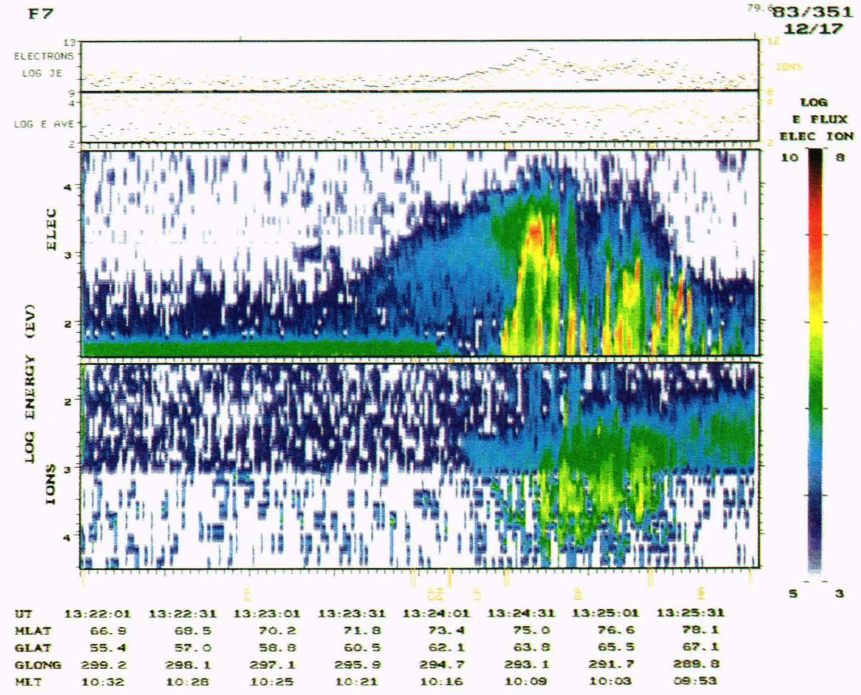

thing) to define a defensible boundary between CPS and BPS, our experience is that when such a boundary clearly exists, the automated identifications are almost invariably reasonable (i.e., a defensible choice, if not necessarily the only choice).

Figure 3C shows a DMSP F7 pass from 5 December 1983 through the dayside oval region, illustrating some of the complex structure that can be observed at these local times. From right to left (equatorward to further poleward) in the spectrogram, one encounters first a region of kiloelectronvolt electron precipitation (CP) (the smooth green cloud). These are electrons originating on the nightside and drifting to the dayside. Further poleward is a region of generally softer electrons with discrete structure; this region is the dayside extension of the oval of discrete aurora (BP). Further poleward is precipitation with characteristics suggestive of the LLBL (LL); next is the region of magnetosheathlike precipita-

tion, the cusp proper (CU); finally, weak precipitation associated with the plasma mantle[8] (MA) lies farthest poleward.

Not all of these regions are observed in each pass, of course. Figure 3D shows a common alternative oval structure observed on 17 December 1983 at only a slightly earlier local time. From left to right (equatorward to poleward), one observes photoelectrons only (PH); both photoelectrons and CPS (CP, PH); a region of CPS without photoelectrons (CP); a region of dayside discrete aurora (BP); and farthest poleward, field lines with precipitation characteristic of the plasma mantle (MA).

FUTURE WORK AND CONCLUSIONS

The most needed improvement is to expand the system to include the dawn-to-dusk satellites (F6 and F8). That will require first constructing an accurate classification scheme for auroral precipitation in those time regions. In making the dayside identifications, we relied on our own work and experience in identifying the source region of particle precipitation. On the nightside, we simply adopted the conventional distinction between diffuse and discrete aurora. At dusk and especially at dawn, however, regions of precipitation exist that do not precisely match any of the nightside or dayside categories; hence, we believe that basic research is required before it is worthwhile to teach a computer to make identifications. The limitation here is not technical but rather our state of knowledge of precipitation and its source regions.

A major technical enhancement that would probably be a significant improvement is a fuller use of context in having the neural network make its decisions. Three likely means of doing this are the following: (1) run the string of one-second decisions through a second neural network, (2) make the neural network much larger so that it can handle simultaneously the full 1500 seconds (by forty input nodes) for a pass, and (3) use a windowing procedure in which the neural network looks at neighboring spectra while making the identification. Progress here requires no new space physics knowledge but only improved computational sophistication.

We conclude by briefly commenting on the advantages and disadvantages of the neural-network–based system as compared with a quantitative-rule–based approach. The primary advantage of a neural network is that it minimizes the work of the space scientist. Either approach would have required the development of a comprehensive classification scheme for identifying each type of precipitation (to know it when you see it). For a quantitative-rule–based system, a great deal of further work would have been necessary to develop and justify a set of explicit rules. For the neural network approach, once the classification effort was completed, the project became a fairly straightforward exercise in applied computer science. The neural network approach undoubtedly minimized the work of the space scientist. An unresolved question, possibly one of philosophy, is whether an explicit and quantitative algorithm for identifying each region represents a deeper level of understanding, as well as more work.

REFERENCES

[1] Hardy, D. A., Schmitt, L. K., Gussenhoven, M. S., Marshall, F. J., Yeh, H. C., et al., *Precipitating Electron and Ion Detectors (SSJ/4) for the Block 5D/Flights 6–10 DMSP Satellites: Calibration and Data Presentation,* Report No. AFGL-TR-84-0317, Air Force Geophysics Laboratory, Hanscom Air Force Base, Mass. (1984).

[2] Winningham, J. D., Yasuhara, F., Akasofu, S.-I., and Heikkila, W. J., "The Latitudinal Morphology of 10-eV to 10-keV Electron Fluxes during Magnetically Quiet and Disturbed Times at the 2100–0300 MLT Sector," *J. Geophys. Res.* **80**, 3148–3171 (1975).

[3] Winningham, J. D., and Heikkila, W. J., "Polar Cap Auroral Electron Fluxes Observed with Isis 1," *J. Geophys. Res.* **79**, 949–957 (1974).

[4] Heikkila, W. J., and Winningham, J. D., "Penetration of Magnetosheath Plasma to Low Altitudes through the Dayside Magnetospheric Cusps," *J. Geophys. Res.* **76**, 883–891 (1971).

[5] Newell, P. T., and Meng, C.-I., The Cusp and the Cleft/LLBL: Low Altitude Identification and Statistical Local Time Variation," *J. Geophys. Res.* **93**, 14,549–14,556 (1988).

[6] Newell, P. T., and Meng, C.-I., "On Quantifying the Distinctions between the Cusp and the Cleft/LLBL," in *Electromagnetic Coupling in the Polar Clefts and Caps,* Sandholt, P. E., and Egeland, A., eds., Kluwer Academic Publishers, Dordrecht, pp. 87–101 (1989).

[7] Lundin, R., "Acceleration/Heating of Plasma on Auroral Field Lines: Preliminary Results from the Viking Satellite," *Ann. Geophys.* **6**, 143–152 (1988).

[8] Newell, P. T., Burke, W. J., Meng, C.-I., Sanchez, E. R., and Greenspan, M. E., "Identification and Observations of the Plasma Mantle at Low Altitude," *J. Geophys Res.* (in press).

[9] McCulloch, W. S., and Pitts, W., "A Logical Calculus of the Ideas Immanent in Nervous Activity," *Bull. Math. Biophys.* **5**, 115–133 (1943).

[10] Hebb, D. O., *The Organization of Behavior*, Wiley, New York (1949).

[11] Rosenblatt, F., *Principles of Neurodynamics*, Spartan Publishers, New York (1962).

[12] Minsky, M., and Papert, S., *Perceptrons,* MIT Press, Cambridge, Mass. (1969).

[13] Rumelhart, D. E., and McClelland, J. L., eds., *Parallel Distributed Processing*, Volume 1, MIT Press, Cambridge, Mass. (1986).

[14] Sejnowski, T. J., and Rosenberg, C. R., "Parallel Networks that Learn To Pronounce English Text," *Complex Sys.* **1**, 145–168 (1987).

[15] Lippman, R. D., "An Introduction to Computing with Neural Nets, *IEEE ASSP Mag.*, 4–22 (1987).

[16] Gorman, R. P., and Sejnowski, T. J., "Analysis of Hidden Units in a Layered Network Trained to Identify Sonar Targets," *Neural Networks* **1**, 75–89 (1988).

[17] Sigillito, V. G., Wing, S., and Hutton, L. V., "Classification of Radar Returns from the Ionosphere Using Neural Networks," *Johns Hopkins APL Tech. Dig.* **10**, 262–266 (1989).

[18] Baker, K. B., and Wing, S., "A New Magnetic Coordinate System for Conjugate Studies at High Latitudes," *J. Geophys. Res.* **94**, 9139–9143 (1989).

[19] Newell, P. T., Wing, S., Meng, C.-I., and Sigillito, V., "The Auroral Oval Position, Structure, and Intensity of Precipitation from 1984 Onwards: An Automated Online Database," *J. Geophys. Res.* (in press).

THE AUTHORS

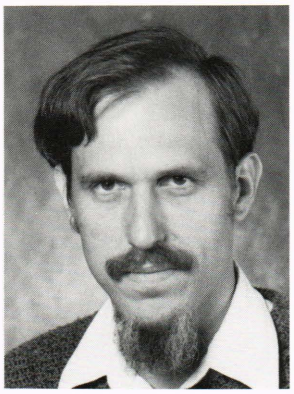

PATRICK T. NEWELL received a B.S. degree from Harvey Mudd College in 1979 and a Ph.D. from the University of California, San Diego, in 1985; both degrees are in physics. His thesis work was based on a sounding rocket experiment conducted over Lima, Peru, to study space plasma physics processes. He joined APL as a postdoctoral researcher in late 1985 and became a member of the Senior Professional Staff in early 1989. Dr. Newell is an author of thirty-four refereed publications (senior author of twenty). His research interest within the specialty of magnetospheric physics lies in the phenomenology of particle precipitation into the ionosphere.

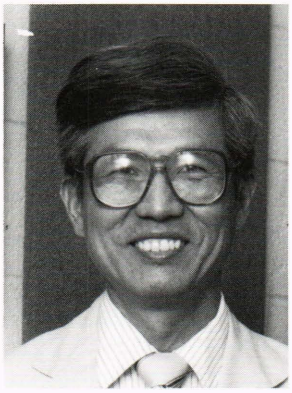

CHING-I. MENG came to the United States in 1963 to study polar geophysical phenomena, such as the aurora and geomagnetic storms, at the Geophysical Institute of the University of Alaska. From 1969 to 1978, he was a research physicist at the Space Sciences Laboratory of the University of California, Berkeley, specializing in magnetospheric physics. Dr. Meng joined APL in 1978 and is involved in the investigation of solar–terrestrial interactions, plasma and field morphology of the magnetosphere, spacecraft charging, and the terrestrial atmosphere. He is also studying the global imaging of the aurora and atmospheric emissions. Dr. Meng became head of APL's Space Sciences Branch in 1990.

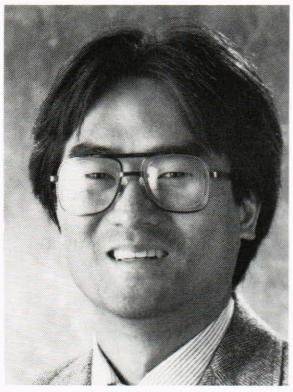

SIMON WING received a B.S. degree in physics in 1984 and an M.S. degree in computer science from The Johns Hopkins University in 1988, where he is now a doctoral student in computer science. From 1984 to 1985, Mr. Wing was with GTE Telenet, where he worked on the PABX project. In 1985, he joined APL's Space Physics Group. Mr. Wing has done extensive software development for the APL high-frequency radar. His interests include neural network research.

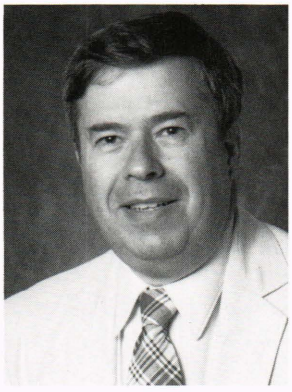

VINCENT SIGILLITO received a B.S. degree in chemical engineering (1958), and M.A. (1962) and Ph.D. (1965) degrees in mathematics from The University of Maryland. He joined APL in 1958 and is supervisor of the Mathematics and Information Science Group in the Milton S. Eisenhower Research Center. Dr. Sigillito is also chairman of the Computer Science Program Committee of the G.W.C. Whiting School of Engineering's Continuing Professional Programs, and has served as interim Associate Dean of the Continuing Professional Programs.