Neural Networks for Oil Spill Detection Using ERS-SAR Data

Fabio Del Frate, Andrea Petrocchi, Juerg Lichtenegger, and Gianna Calabresi

Abstract—A neural network approach for semi-automatic detection of oil spills in European remote sensing satellite-synthetic aperture radar (ERS-SAR) imagery is presented. The network input is a vector containing the values of a set of features characterizing an oil spill candidate. The classification performance of the algorithm has been evaluated on a data set containing verified examples of oil spill and look-alike. A direct analysis of the information content of the calculated features has been also carried out through an extended pruning procedure of the net.

Index Terms—ERS-synthetic aperture radar (SAR), neural networks, oil spill detection.

I. INTRODUCTION

THE DEVELOPMENT of spaceborne synthetic aperture radar (SAR) systems, in particular the advent of the European remote sensing satellite (ERS) missions, have improved the possibilities for the detection of oil spills, seriously affecting the marine ecosystem, allowing a more rigorous and cost-effective monitoring. The presence of an oil film on the sea surface damps the small waves due to the increased viscosity of the top layer and drastically reduces the measured backscattering energy, resulting in darker areas in SAR imagery [1], [2]. However, careful interpretation is required because dark areas might also be caused by locally low winds or by natural sea slicks. To avoid false alarms, experienced image interpreters or well tuned classification algorithms are employed. Indeed, oil spills show a larger discontinuity effect with respect to the background, mainly because of their high viscosity.

Based on such a rationale, a possible operational procedure could be as follows:

1) selection of an area in the image containing a dark object;
2) computation of physical and geometrical features characterizing the object;
3) classification of the object into oil spill or look-alike, based on the calculated features.

In this paper, we considered an operational scenario like the one illustrated in Fig. 1. We assumed that an operator, by visual inspection, performs the first step, while the two last steps are executed automatically. Even though the attention of this work
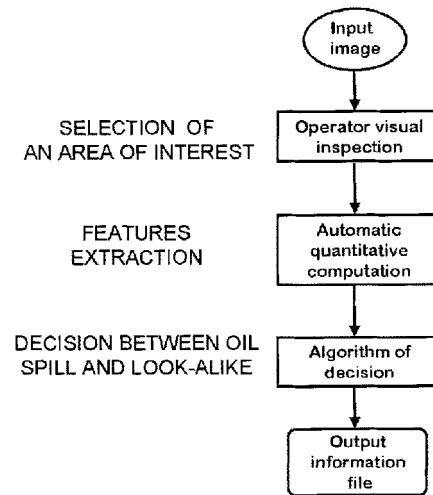
Fig. 1. Block-diagram showing a possible operational scenario for oil spill detection.

is focused on the classification algorithm, a description of the general aspects for the features extraction will be also given.

Mostly classification algorithms for oil spill detection rely on Bayesian or other statistically-based decisions. The drawback of these methods is the complex process to develop classification rules, due to the many nonlinear and poorly understood factors involved. These kinds of difficulties can be overcome by considering a neural network approach.

The use of neural networks in remote sensing has often been found effective, since they can simultaneously handle nonlinear mapping of a multidimensional input space onto the output one and cope with complex statistical behavior [3]. Neural networks, conversely, from statistically-based classifiers, do not require an explicitly well defined relationship between the input and the output vectors, as they determine their own input-output discriminant relations directly from a set of training data, used to draw the decision boundaries [4].

Our study investigates the potentiality of a neural network approach to provide a classification algorithm only using ERS-SAR data without auxiliary information. This could be used to support and make the decision-taking process of image analysts more reliable. Moreover, if an appropriate filter that selects dark objects in the image is designed, the algorithm can be considered for a fully automatic approach.

The input of the net consists of a set of features providing information about an oil spill candidate, and the output gives the probability for the candidate to be a real oil spill. An optimization of the neural net has also been carried out by using pruning
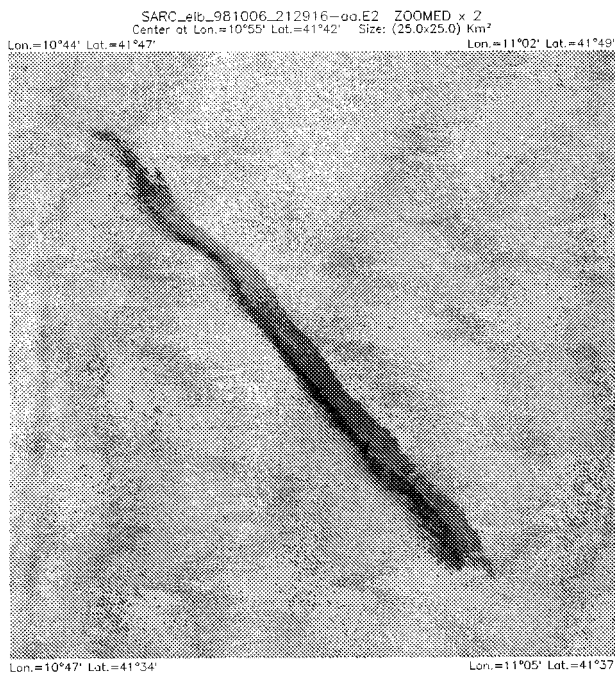
Fig. 2. Image of a verified oil spill taken on June 10, 1998 outside Elba Island (Italy).

Fig. 3. Image of a verified natural film taken on July 5, 1998, near the coast near Palermo, Italy.

algorithms which eliminate its ineffective connections. A discussion on the performance of the net and on the information content of the considered individual features characterizing the input vector is also conducted.

II. DATA SET

The nominal spatial resolution of the ERS data is 25 m × 25 m, but for oil spill detection, low resolution images of 100 m × 100 m are more than sufficient, and the data amount of each image significantly reduced. For this reason, fast delivery images produced at Fucino Station, Italy, have been considered in our work. These products are common binary files of 1050 Kb reduced from a multilook SAR image by resampling (5 × 6 pixels block averaging).

It is to be noted that these products can be generated immediately after data acquisition and hence allow us to set up a system that could support Coastal Guards for their activity of marine pollution control.

For promotion and demonstration purposes, over 600 images have been taken in the years 1997 and 1998 over various areas of the Mediterranean Sea. Out of this archive, each image containing at least one occurrence of look-alike or of oil spill, has been selected and calibrated so that its pixel intensity values were representing the measured backscattering coefficients. Images with examples of oil spill and of look-alike are shown in Figs. 2 and 3, respectively. Many of these occurrences were reported to the authorities and verified by the Coast Guards of Italy and France.

After the radiometric range correction and the georeferencing of each image, a dedicated oil spill processing and analysis tool is applied, based on edge detection. The procedure starts with the definition of a region of interest by the user. The tool analyzes the overall backscattering of the region and in particular
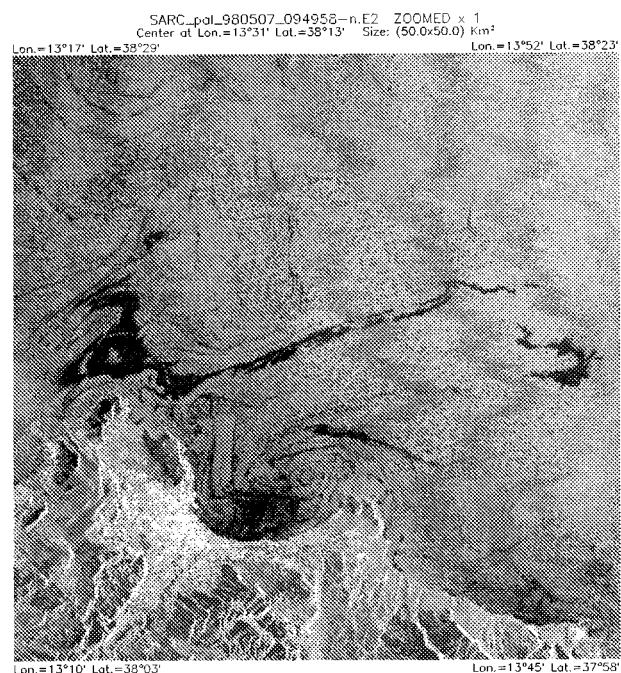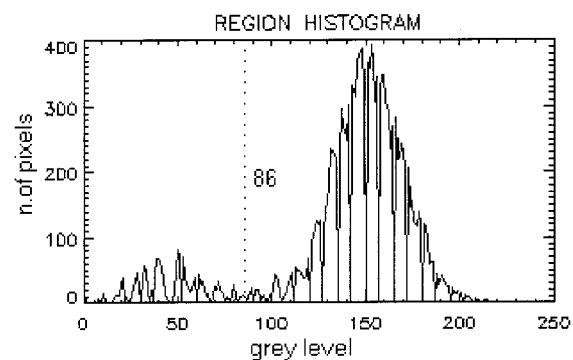

Fig. 4. Histogram where the number of pixels in the region with the same grey level is reported. The dashed line corresponds to the local minimum between the two peaks and indicates the threshold value needed for the edge detection of a oil spill candidate.

produces a histogram. The typical shape of such a histogram is shown in Fig. 4. We see that the histogram contains two peaks, the lower is located around the mean backscattering value of the dark object, the taller around the mean value of the background. The local minimum value between the two peaks is stored and is the one used for image fragmentation. To this purpose, the darkest pixel in the region is selected as a starting point and then the region around this pixel grows, defining an edge that corresponds to the border of the oil. The region grows until the neighboring pixels have a value greater of the threshold value given by the local minimum previously calculated (Fig. 4). Several routines have been implemented dealing with anomalies from the described general case. Subsequently, the human operator can either accept the result suggested by the automatic procedure or reject it and produce a new edge detection by changing manually the threshold. In this way, the oil spill analysis tool is made flexible and adaptive to a variety of situations. Once the
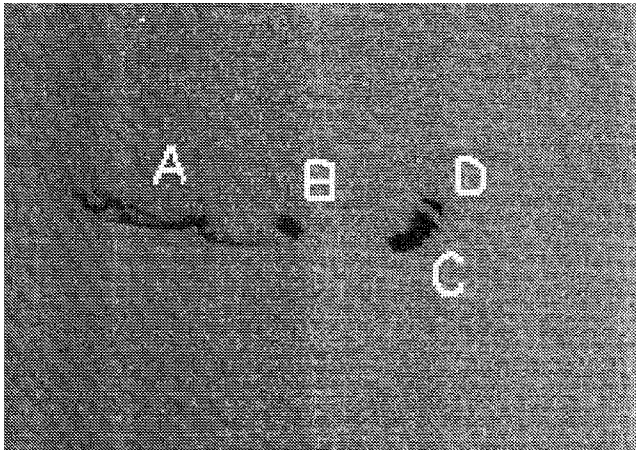
Fig. 5. Fragmented oil spill. Image taken on August 6, 1998 outside Sicily Island, Italy.

border of the dark object is accepted, a mask is generated and a number of morphological and physical parameters are computed. If more than one dark object is present in the selected area, the described process can be extended to a neighboring object so that oil spills consisting of multiple releases can be also analyzed. In these cases, the features are first calculated for each element and then elaborated to describe one single representative object. An example of such a situation is reported in Fig. 5, where the dark objects A, B, C, and D can be reasonably assumed to originate from a single (likely) oil spill event.

At this point, the features extracted from the dark object are input to a classifier that estimates its probability of being an oil slick. These features consider the geometry of the dark object in terms of its extension and of its shape, as well as the physical behavior in terms of the characteristics of the backscattering intensity of the pixels belonging to the object, to the background, and to the area around the border. A more detailed description of the features is the following.

1) Area (A). Area (in km^2) of the object.
2) Perimeter (P). Length (in km) of the border of the object.
3) Complexity (C). Complexity is defined as follows:

$$C = \frac{P}{2\sqrt{\pi A}}. \qquad (1)$$

This feature will generally take a small numerical value for regions with simple geometry and larger values for complex geometrical regions.

4) Spreading (S). This feature is derived from the principal component analysis [5] of the vectors whose components are the coordinates of the pixels belonging to the object. If λ_1 and λ_2 are the two eigenvalues associated with the computed covariance matrix and $\lambda_1 > \lambda_2$, the spreading value S is computed using this expression

$$S = \frac{100\lambda_2}{\lambda_1 + \lambda_2}. \qquad (2)$$

S will be low for long and thin objects and high for objects closer to a circular shape.

TABLE I
MAIN STATISTICAL PARAMETERS DESCRIBING THE FEATURES EXTRACTED FROM THE DATA SET

Feature	OIL SPILL				LOOK ALIKE			
	Min	Max	Mean	St. Dev.	Min	Max	Mean	St. Dev.
A (km^2)	0.4	40.6	6.4	7.5	1.1	115.6	13.3	17.0
P (km)	4.2	117.6	28.1	22.1	7.1	396.4	52.4	57.2
C	1.1	6.8	3.2	1.1	1.1	10.4	3.9	1.7
S	0.1	40.8	4.2	7.4	0.1	45.2	11.8	11.4
OSd (dB)	0.8	3.8	1.7	0.6	0.9	3.2	2.0	0.6
BSd (dB)	0.8	3.0	1.1	0.4	0.9	2.3	1.5	0.4
ConMax (dB)	3.2	15.7	9.2	2.6	2.6	14.9	10.8	2.2
ConMe (dB)	-0.4	10.9	4.8	1.9	-0.4	9.3	5.3	1.7
GMax (dB)	2.8	15.5	7.2	2.2	3.6	16.8	8.5	2.6
GMe (dB)	0.0	6.5	3.0	1.1	0.0	5.2	2.7	1.0
GSd (dB)	0.8	2.7	1.4	0.4	0.6	2.6	1.5	0.5

5) Object standard deviation (OSd). Standard deviation (in dB) of the intensity values of the pixels belonging to the oil spill candidate.
6) Background standard deviation (BSd). Standard deviation (in dB) of the intensity values of the pixels belonging to the region of interest, selected by the user, surrounding the object.
7) Max contrast (ConMax). Difference (in dB) between the background mean value and the lowest value inside the object.
8) Mean contrast (ConMe). Difference (in dB) between the background mean value and the object mean value.
9) Max gradient (GMax). Maximum value (in dB) of border gradient.
10) Mean gradient (GMe). Mean border gradient (in dB).
11) Gradient standard deviation (GSd). Standard deviation (in dB) of the border gradient values.

The described set of features has been calculated for 139 dark objects, 71 oil spills, and 68 look-alikes. For a number of examples, ground truth was available. Otherwise, the discrimination was based on the independent judgment of experienced image analysts. In Table I, some statistical parameters of the features extracted from our data set are reported. The table shows that oil spills have less complex and more thin shapes, while in case of the look-alikes, the values of the backscattering either in the object or in the surroundings are more dispersed. Also, oil spills show mean value of the gradient along the border higher than look-alikes. Look-alikes are generally larger than oil spills.

III. NEURAL ALGORITHM

An artificial neural network (NN) may be viewed as a mathematical model composed of many nonlinear computational elements, named neurons operating in parallel and massively connected by links characterized by different weights. A single neuron computes the sum of its inputs, adds a bias term, and drives the result through a generally nonlinear activation function to produce a single output termed the activation level of that neuron. NN models are mainly specified by the net topology, neuron characteristics, and training or learning rules.

The term topology refers to the structure of the network as a whole, specifying how its input, output, and hidden units are interconnected [6]. For this study, multilayer perceptrons (MLP) have been considered, which have been found to have the best suited topology for classification and inversion problems [7].
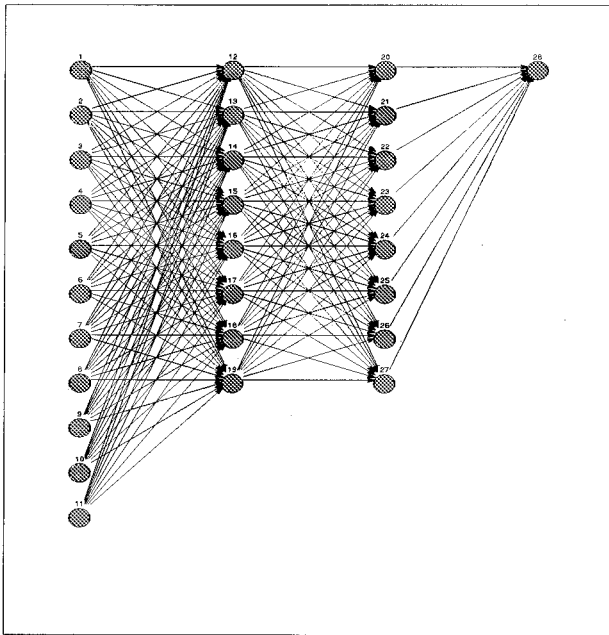
Fig. 6. Feedforward topology for a neural network (NN).

These are feedforward networks, where the input flows only in one direction to the output, and each neuron of a layer is connected to all neurons of the successive layer but has no feedback to neurons in the previous layers. Two layers of weights have been used in our case. The described topology is reported in Fig. 6.

The individual neuron is the elemental building block of each layer, and it is mainly characterized by its activation function. The most common activation function is the nonlinear sigmoid function, also used in our simulations, defined as follows:

$$f(x) = \frac{1}{1 + e^{-x}}. \tag{3}$$

Such an activation function yields values in the range [0,1].

An MLP is designed to approximate an unknown input–output relation by determining the weight or strength of each connection via learning rules. These rules indicate how to pursue minimization of the error function measuring the quality of the network's approximation on the domain covered by a training set (i.e., a set of input–output examples). In our case, the input vector contains the measured features, the output vector contains the response classification provided either by the ground truth, or by image interpretation experts. A preprocessing of the input variables is commonly performed. In particular, this is useful if different input variables have typical values that differ significantly and do not reflect the relative importance in determining the required outputs. For these reasons we scaled the input values so that they were between 0.01–0.99.

The net was trained using the back propagation algorithm, which uses a gradient search technique and iteratively adjusts the weight coefficients in the network to minimize an error function equal to the mean square difference between the desired and the actual net output. Iterating the procedure, the learning process was stopped when no more significant variations in the overall error were observed.

The neural network simulator (SNNS) developed at the University of Stuttgart, Stuttgart, Germany, [8] formed the basic software for the classification algorithm implementation and proved to be a high level, flexible, and reliable software package.

IV. RESULTS

Several attempts have been initially made to properly select the number of units to be considered in the hidden layers. The topology 11-8-8-1 has been finally chosen for its good performance as well in terms of classification accuracy as well of training time. In fact, the number of about 15 000 training cycles was sufficient to get the network learned.

The classification results are very promising as, once trained, the net is able to correctly classify all examples. The neural responses are generally not ambiguous. In fact, the outputs tend to assume the extreme values in the allowed range [0,1]. This means that with the considered data set and with the considered learning strategy the net succeeded in determining a precise mapping function. The function returns 0 if a dark spot in the image is an oil spill, 1 if it is a look-alike.

A step forward in the evaluation of the neural net performance consists in the analysis of its behavior once it is tested on a set of new examples, not belonging to the training set. In order to avoid the long time consuming work for the collection of a new set of data, we decided to adopt the leave-one-out method [9]. According to this procedure, if N is the number of the available data points, a net is trained using $N-1$ points and tested on the remaining one. This process is repeated for each of the N possible choices, and the results can give a first estimation on the generalizing capabilities of the net. In our case, where $N = 139$, it has been found that the 18% of real oil spills has been misclassified as look-alike, and 10% of look-alikes have been misclassified as oil spill, with an overall rate of misclassified pixels of 14%. So there are some situations in which the single data point does not respect the internal rules that the net has learned from the other points. This can be explained with the heterogeneity of the data set, which includes several examples with very different geometric and physical characteristics.

The good performance of the neural algorithm encouraged a second phase where an optimization of the net from the point of view of the number of its adaptive parameters (units and connections) has been carried out by using a pruning procedure. According to this, a network is examined to assess the relative importance of its weights, and the least important ones are deleted. Typically, this is followed by some further training of the pruned network, and the procedure of pruning and training may be repeated for several cycles. Clearly, there are various choices to be made concerning how much training is applied at each stage, what fraction of the weights is pruned, and so on. In our case, every time a weight was removed, we trained the new net until, as in the case of the initial training, the overall error value was approaching a value of convergence, and, since we started with a net committing no errors, we continued with the pruning procedure until we realized that new removals involved errors in
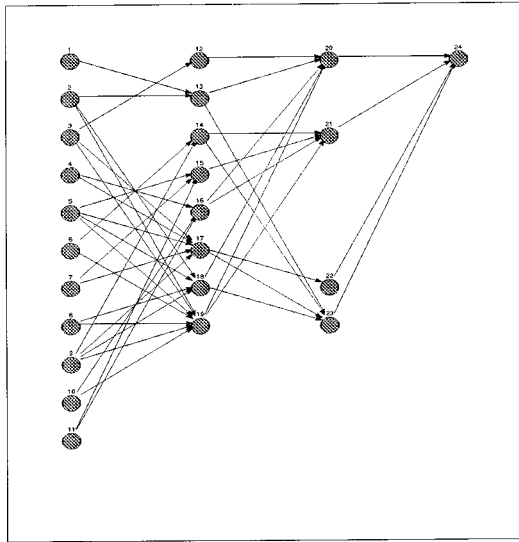
Fig. 7.　NN after the pruning procedure.

TABLE II
CLASSIFICATION OF THE FEATURES SIGNIFIGANCE ACCORDING TO THE
ORDER OF REMOVAL OF THE INPUT COMPONENTS BY APPLYING
THE PRUNING ALGORTIHM

Level of Significance	Features
1	GSd, GMe GMax, BSd
2	C, OSd ConMe, S
3	A, P ConMax

the classification task. The most important consideration, however, is how to decide which weights should be removed. To do this, we need some measure of the relative importance, or saliency, of different weights. The results have been obtained by applying the simple concept that small weights are less important than large weights and by using the magnitude of a weight value as a measure of its importance. At the end of the pruning process, we obtained a new net that, though in terms of classification accuracy showed the same performance of the initial one, was characterized by a much smaller number of connections: 45 versus 160. This net is illustrated in Fig. 7.

We finally tried to examine which features, in the chosen context, were containing less information and, conversely, which were the features which might be crucial for the success of classification in the oil slick detection on SAR imagery. To this end, we prolonged the pruning procedure to the input layer, until ten of the eleven components of the input vector were removed (it has to be recalled that an input or hidden unit is removed when it has lost all its connections). The result of this analysis is synthetically reported in Table II, where the features are classified through three levels of significance, deduced from the order of removal of the corresponding input units. We see that almost

all features belonging to the first class, i.e., containing the most useful information, are those mainly involved in quantifying the border discontinuity effect, so basically, those containing information on the gradient of the backscattering value when we pass from the background into the spill. A fourth crucial quantity seems to be the background standard deviation, which is actually one of the parameters most affected by wind effects and is generally high for natural sea slicks. In fact, natural sea slicks tend to be more sparsely distributed than oil spills on the sea surface. This can provoke the main film (i.e., the analyzed dark object) to be accompanied by other minor damping effects in the surrounding area (see also Fig. 3).

In the second level, shape features, together with a measure of the homogeneity of the backscattering intensity inside the object, confirm to give a significant contribution of information. Finally, features measuring the extension of the dark object are found to be reasonable in the third class.

V. CONCLUSIONS

The potentialities of NN algorithms for the detection of oil spills in ERS-SAR imagery have been studied in this paper. The data necessary to train the net has been generated using routines that retrieve a set of specific features describing a dark area in a ERS-SAR image over the sea. The neural net could correctly discriminate over a set of independent examples between oil spills and look-alikes with a largely acceptable rate of success. It has been shown that the net can be made much simpler than its initial and standard feedforward configuration. More than 70% of the connections of the initial topology could be removed by means of a pruning algorithm without loss of accuracy of the classification. Extending the pruning to the input layer, the significance of the individual feature could be assessed.

It must be also pointed out that additional features taking into account local atmospheric conditions such as wind speed might still improve the oil spill detection performance.

REFERENCES

[1] W. Alpers and H. Hühnerfuss, "Radar signatures of oil films floating on the sea surface and the Marangoni effect," *J. Geophys. Res.*, vol. 93, pp. 3642–3648, Apr. 1988.
[2] P. Trivero, B. Fiscella, F. Gomez, and P. Pavese, "SAR detection and characterization of sea surface slicks," *Int. J. Remote Sensing*, vol. 19, pp. 543–548, 1998.
[3] M. S. Dawson, "Applications of electromagnetic scattering models to parameter retrieval and classification," in *A. Fung Microwave Scattering and Emission Models and their Application*. Norwood, MA: Artech House, 1994.
[4] D. E. Rumelhart, G. E. Hinton, and R. J. Williams, "Learning internal representations by error propagation," in *Parallel Distributed Processing*, D. E. Rumelhart and J. L. McClelland, Eds. Cambridge: MIT Press, 1986, pp. 318–362.
[5] I. T. Jollife, *Principal Component Analysis*. New York: Springer-Verlag, 1986.
[6] R. P. Lippmann, "An introduction to computing with neural nets," *IEEE ASSP Mag.*, pp. 4–22, 1987.
[7] S.-Y. Hsu, T. Masters, M. Olson, M. Tenorio, and T. Grogan, "Comparative analysis of five neural network models," *Remote Sensing Rev.*, vol. 6, pp. 319–329, 1992.
[8] A. Zell, *et al.*, "SNNS Stuttgart neural network simulator user manual," Inst. Parallel Distributed High Performance Syst., Univ. Stuttgart, Stuttgart, Germany, 1995.
[9] C. M. Bishop, *Neural Networks for Pattern Recognition*. Oxford, U.K.: Oxford Univ. Press, 1995, pp. 374–375.
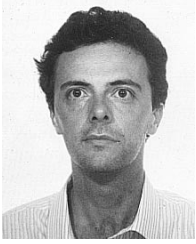
Fabio Del Frate received the Laurea degree in electronic engineering in 1992 and the Ph.D. degree in computer science in 1997, both from the University of Rome Tor Vergata, Rome, Italy.

From September 1995 to June 1996, he was a Visiting Scientist with the Massachusetts Institute of Technology, Cambridge, and from 1998 to 1999, he was with the ESRIN Center, European Space Agency, Frascati, Italy, as a Research Fellow. Currently he is with the University of Rome Tor Vergata, involved in programs of teaching and research on electromagnetics and remote sensing. Since 1992, his research activity has been mainly dedicated to inversion problems in active and passive remote sensing.

Juerg Lichtenegger received the M.S. and Ph.D. degrees from the Department of Geography, University of Zürich, Zürich, Switzerland, in 1975 and 1980, respectively.

He has been an Engineer in the User Service with the European Space Agency, Ispra, Italy, since 1982, working mainly with Landsat TM and coordinating European SAR aircraft experiments. With the launch of ERS, he got involved in project monitoring and in application development and demonstrations for the SAR. This activity also includes supervision of trainees and students and lecturing in training courses, in-house and abroad. In this position, he is the co-author in several papers regarding image processing with respect to SAR multitemporal image analysis in general and for environmental monitoring and spill detection in particular.

Andrea Petrocchi received the degree in astrophysics and space physics from the University of Rome "La Sapienza," Rome, Italy, in 1992.

With the Planetological Deptartment, CNR Institute of Astrophysics, Rome, Italy, he analyzed satellite data from stars and galaxies and later was involved in producing thematic maps of the surface and the atmosphere of planet Mars. He has been with the European Space Agency since 1994, contributing for two years to the quality control of ERS-1 data and participating in the "commissioning phase" after the launch of ERS-2. Now he is developing data applications and following projects mainly related to SAR data. His main activity today is focused on oil spill detection by ERS-SAR.

Gianna Calabresi received the Ph.D. degree in political sciences from Rome University "La Sapienza," Rome Italy, in 1976.

Following previous working experiences with an oil company first, she was with Telespazio, Italy, in 1980. She then joined the European Space Agency (ESA), Frascati, Italy, as an Engineer in the User Service of the Earthnet Programme Office. Her role was, firstly, user and ground station interface for data products derived from Landsat and similar EO missions, and included the coordination of National Points of Contact within Europe, institutionally in charge of remote sensing activities in their countries. Later, she was and is presently involved in ERS SAR application projects, new user development, promotion of pre-operational use of SAR-derived information in oil spill detection, floods etc., and support to ESA's subsidized Announcement of Opportunities and Data User projects.