

Optimum guard zone for self-supervised learning

S.K. Pal, M.Tech., Ph.D., Mem.I.E.E.E.

Indexing terms: Algorithms, Fuzzy sets, Pattern recognition, Vowel recognition

Abstract: A self-supervised learning algorithm using fuzzy set and the concept of guard zones around the class representative vectors is presented and demonstrated for vowel recognition. An optimum guard zone having the best match with the fully supervised performance is determined. Results are also compared with that of nonsupervised case for various orders of input patterns.

1 Introduction

An adaptive pattern recognition system can be viewed as a learning machine which improves the system's performance by acquiring necessary information for decision during the system's operation. In a supervised system, the machine in general, requires an extra source of knowledge, usually of a higher order, for correcting the decision taken by the classifier. Bayesian estimation and stochastic approximation [1, 2] can be used for supervised learning to learn unknown parameters successively in a given form of distribution of each class. In a strictly nonsupervised adaptive system, these parameters are updated solely on the basis of the decision of the classifier. The convergence of the system to an optimal set of class representative parameters may thus be seriously affected by incorrect decisions of the classifier [1]. Again, when an extra source of knowledge on which the supervisory programme could be based is not readily available, the performance of the system becomes completely unpredictable. In the present paper, a system of self supervision based on inherent properties of the class distribution of features is proposed. For all practical problems, the distribution of the members of a class in the feature space has a central tendency and it may be assumed that probability of misclassification near these central tendencies is substantially low. Thus, one can construct a region around this central tendency of a class, so that an unrestricted updating procedure for the samples falling only in this region would assist the convergence of the system significantly. Such a region, defined as a 'guard zone', forms the basis of a supervisory programme which needs only to check whether the classified input is within the guard zone for the purpose of inhibition of the updating programme. A mathematical formulation of such guard zones would require a thorough knowledge of the distribution function of the features for each class. When these functions are not precisely known, the size of the guard zones has to be experimentally determined. The various guard zones, the semiaxes of which are the $1/\lambda$ th ($\lambda^2 = 0.5, 1, 2, 4, 6, 8$) part of the corresponding standard deviations are therefore considered around the selected (estimated) representative vectors for obtaining the optimum one.

The effectiveness of the algorithm with a classifier based on the concept of fuzzy set [3] is demonstrated on the different sequences of a set of about 900 vowel sounds when the first three format frequencies are considered as input features.

2 Fuzzy sets and classification algorithm

A fuzzy set A with its finite number of supports x_1, x_2, \dots, x_n in the universe of discourse U is defined as

$$A = \{\mu_A(x_i)/x_i\} \quad i = 1, 2, \dots, n \quad (1)$$

where the membership function $\mu_A(x_i)$ having positive values

in the interval $[0, 1]$ denotes the degree to which an event x_i may be a member of A . This characteristic function can be viewed as a weighting coefficient which reflects the ambiguity (fuzziness) in A .

Similarly, the property p defined on an event x_i is a function $p(x_i)$ which can have values only in the interval $[0, 1]$. A set of these functions, which assigns the degree of possessing some property p by the event x_i constitutes what is called a property set [4].

Now we develop a multicategory classifier on the basis of the property set where the input and output (decision) are deterministic, but the process of classification is fuzzy. Let

$$X = [p_1/x_1, p_2/x_2, \dots, p_n/x_n, \dots, p_N/x_N]$$

be an unknown pattern in an N -dimensional vector space containing m pattern classes to be recognised. p_n/x_n denotes the degree of possessing the n th fuzzy property p_n by the n th measurement x_n of the pattern X and has values between zero and one. The decision of the classifier is based on the magnitude of the fuzzy similarity vector $S_j^{(l)}(X) = \{s_{nj}^{(l)}\}$ of X with respect to the l th prototype in the j th class, where

$$s_{nj}^{(l)} = [1 + |p_{nj}^{(l)} - p_n|/(p_{nj}^{(l)}\sigma_{nj}^{(l)})]^{-1} \quad (2)$$

$$p_n = [1 + |\hat{x}_n - x_n|/F_d]^{-F_e} \quad (3)$$

$$p_{nj}^{(l)} = [1 + |\hat{x}_n - \bar{x}_{nj}^{(l)}|/F_d]^{-F_e} \quad (4)$$

$$\hat{x}_n = \max_j \max_l \{\bar{x}_{nj}^{(l)}\} \quad (5)$$

for $n = 1, 2, \dots, N; j = 1, 2, \dots, m$ and $l = 1, 2, \dots, h_j$.

$s_{nj}^{(l)}$ denotes the grade of similarity between the n th property of X and that of the l th prototype in the j th class. $p_{nj}^{(l)}$ denotes the degree to which property p_n is possessed by the l th prototype in the j th class. \hat{x}_n is the n th reference constant. h_j is the number of prototypes in the j th class. $\bar{x}_{nj}^{(l)}$ and $\sigma_{nj}^{(l)}$ correspond to the l th prototype and represent the mean and standard deviation corresponding to the n th component in the j th class. $(1/\sigma_{nj}^{(l)})$ is used as a weighting coefficient in measuring similarity. Positive constants F_e and F_d are the exponential and denominational fuzzifiers, respectively, which play the role of creating different amount of fuzziness in the property set [5]. It is to be noted here that the property p_n (eqn. 3) is defined using the π -function [6] which represents the compatibility function corresponding to a fuzzy set ' x_n is \hat{x}_n '.

The pattern X is then decided to be from the k th class if

$$|S_k(X)| = \max_j |S_j(X)| \quad (6)$$

where

$$|S_j(X)| = \max_l |S_j^{(l)}(X)| \quad (7)$$

for $j = k = 1, 2, \dots, m$.

3 Self-supervision algorithm

The 'decision parameter of the supervisor' (*DPS*) for the *j*th class is defined as

$$(DPS)_j = \sum_n [(x_n - \bar{x}_{nj}^{(l)})/\sigma_{nj}^{(l)}]^2 \quad (8)$$

$$\sigma_{nj}^{(l)} = \sigma_{nj}^{(l)}/\lambda \quad (9)$$

for $n = 1, 2, \dots, N; l = 1, 2, \dots, h_j$ and $j = 1, 2, \dots, m$.

λ (a positive constant) is termed as 'zone controlling parameter' which controls the dimension $\sigma_{nj}^{(l)}$ of the guard zone in a class. It is to be mentioned here that this decision parameter would lead to ellipsoidal shapes of the guard zones. Since the system uses the inherent properties of the distribution of the same parameters as used by the classifier itself, it may be called a 'self-supervisory system'.

The supervisor then accepts the decision made by the classifier that *X* is from the *k*th class, only if

$$(DPS)_k \leq 1 \quad (10)$$

and the parameters $\bar{x}_{nj}^{(l)}$, \hat{x}_n , $\sigma_{nj}^{(l)}$ and $p_{nj}^{(l)}$ for the *k*th class are then correspondingly updated for that input sample *X*. Otherwise, the decision is thought to be doubtful and no other alteration of these class parameters is made.

4 Iterative algorithm for parameter estimation

In general, the input events which are to be classified are in a somewhat randomly mixed sequence. These samples after being classified become members of certain classes and modify the centres and the weight vectors of them.

Let $\bar{x}_{nj(t)}$, $\sigma_{nj(t)}^2$ and $p_{nj(t)}$ represent the mean, variance and fuzzy property *p*, respectively, along the *n*th co-ordinate axis of the *j*th class and $\hat{x}_n(t)$ be the *n*th reference constant, estimated by the first *t* samples. Then, after the addition of another sample $x_{n(t+1)}$, these parameters would be adjusted as follows:

$$\bar{x}_{nj(t+1)} = \frac{t}{t+1} \bar{x}_{nj(t)} + \frac{1}{t+1} x_{n(t+1)} \quad (11a)$$

$$\sigma_{nj(t+1)}^2 = \frac{1}{t+1} C_{nj(t+1)} - \bar{x}_{nj(t+1)}^2 \quad (11b)$$

$$C_{nj(t+1)} = C_{nj(t)} + x_{nj(t+1)}^2 \quad (11c)$$

$$C_{nj(t)} = \sum_i x_{ij}^2 \quad i = 1, 2, \dots, t \quad (11d)$$

$$\hat{x}_n(t+1) = \max_j \bar{x}_{nj(t+1)} \quad (11e)$$

$$p_{nj(t+1)} = [1 + |\hat{x}_n(t+1) - \bar{x}_{nj(t+1)}|/F_d]^{-F_e} \quad (11f)$$

For simplicity, we have omitted *l* (superscript) of these parameters. It is to be noted here that only the parameters $\bar{x}_{nj(t+1)}$ and $C_{nj(t+1)}$ (eqns. 11a and c) are directly modified by the new input $x_{n(t+1)}$. The others follow from these two.

5 Method of recognition

Fig. 1 shows the block diagram of a self-supervised recognition system. The model uses a classifier based on 'fuzzy property set' which measures the similarity between the different representative vectors and the input vector and then assigns the input to the class for which the representative vectors show maximum similarity. To study the adaptive capability of the system in recognising a pattern, the initial values of these parameters are deliberately chosen to be different ($\langle \bar{x}_{nj(t)} \rangle$, $\langle \sigma_{nj(t)} \rangle$, $\langle \hat{x}_n(t) \rangle$, $\langle p_{nj(t)} \rangle$) from their true values. $\langle \cdot \rangle$ denotes the estimated value. Of these four parameters, as mentioned in the Section 4, we need primarily $\langle \bar{x}_{nj}^{(l)} \rangle$ and $\langle \sigma_{nj}^{(l)} \rangle$ to be estimated from some sets of training samples. Others are being automatically derived from those estimates.

After the classification of *X*, it is the task of the supervisor to judge whether the sample *X* is within the specified guard zone as defined around $\langle \bar{x}_{nj}^{(l)} \rangle$. If it does, the decision of the classifier that *X* is from *k*th class is accepted by the supervisor and the parameters of *k*th class are updated by *X*. Otherwise, there will be no alteration of the class parameters before the next input.

In fully supervised learning, the decision of the classifier is verified by an external supervisor and the class parameters are altered only if the classification is found to be correct.

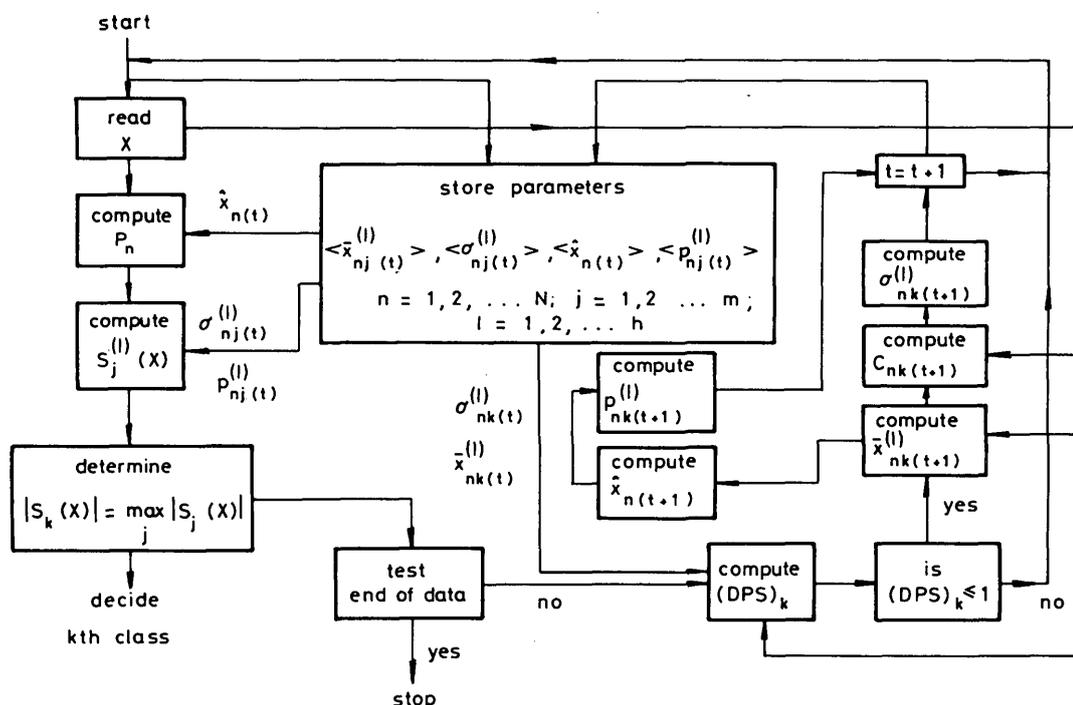


Fig. 1 Block diagram of self-supervised recognition system

For nonsupervised case, the decision of the classifier is considered to be final and the parameters of the recognised classes are correspondingly modified.

6 Implementation of vowel recognition

The previously mentioned algorithm was implemented on a set of 871 Telugu (an important Indian language) vowel sounds [7-10] uttered by three speakers in the age group of 30 to 35 years. The first three vowel formant frequencies were considered as recognition features to classify ten vowel classes (∂ , a:, i, i:, e, e:, u, u:, o and o:) including long and short categories. Since the short and long categories of a vowel differ only in duration, these ten vowels were then divided into six groups (∂ , a:, I, E, U and O) which differ only in phonetic features.

The set of data for each of the vowel classes has been found to follow the normal distribution [8]. Therefore, the assumption (as mentioned in Section 1) that 'the probability of misclassification of the input patterns falling within the guard zone constructed around the central tendency of a class distribution is substantially low' is well justified here. An unrestricted updating programme for those samples would thus assist the convergence of the system significantly. Of course, the convergency of this adaptive system (as mentioned in the following text) is also experimentally verified [9].

Now, we are interested here in studying the adaptive efficiency of the system in recognising vowel sounds with the

nonappropriate prototype vectors representing the classes. Similar investigations have also been reported [9] which used nonadaptive, fully supervised and nonsupervised procedures where the prototype points and corresponding weighting coefficients of a specified class were obtained from five utterances of one of the three speakers selected randomly from each of the classes. Such an initial incorrect set of class representative points was found as the process of classification continued, to approach gradually the respective true mean values, demonstrating the convergence property of the learning algorithm. Once the optimum size of the training set (usually 16 to 20 samples per class) [8-10] is obtained by the classifier, further increase in the size of the set does not improve the system's adaptivity, and hence the performance, significantly. The initial class representative vectors in this experiment were chosen just outside the boundary of an ellipsoid having the three axes equal to the respective standard deviations of the features and mean of the class as the centre. The standard deviations for providing weighting coefficients corresponding to those representative points were obtained from a set of ten training samples selected randomly from each of the classes. Although the shorter and longer types of vowels I, E, U and O are treated similarly, they were given individual reference and weighting vectors. Thus, in our experiment, $m = 6$, $N = 3$, $h = 1$ for ∂ and a:, and $h = 2$ for I, E, U and O. F_e and F_d were considered to be 0.5 and 40000, respectively. These values of the fuzzifiers were found to be optimum [5, 8, 11] with respect to recognition score.

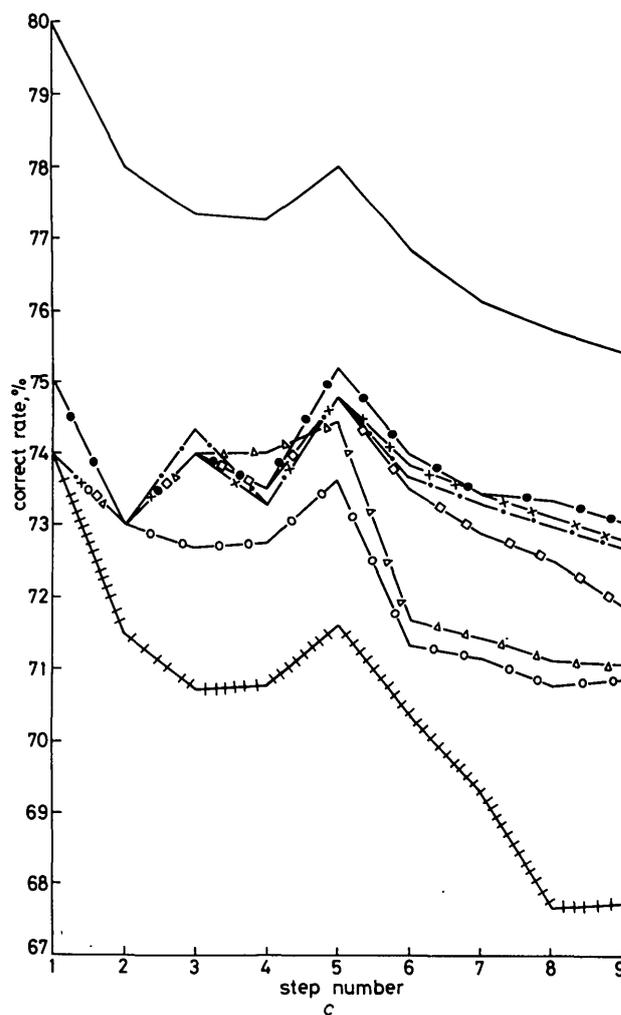
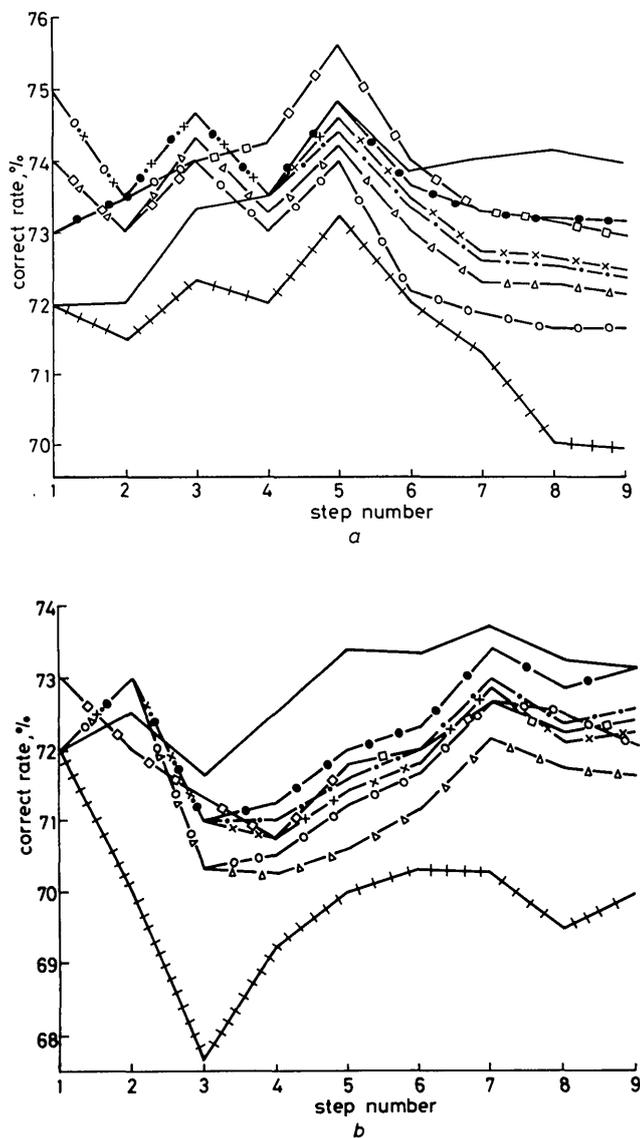


Fig. 2 System performance curves

— fully supervised, ++++ nonsupervised
 □-□-□ $\lambda^2 = 0.5$, ○-○-○ $\lambda^2 = 1$
 △-△-△ $\lambda^2 = 2$, ●-●-● $\lambda^2 = 4$
 x-x-x $\lambda^2 = 6$, - - - $\lambda^2 = 8$

7 Experimental results

Since the performance of an adaptive system depends much on the sequence of incoming samples [1], the experiment was repeated ten times for different orders of appearance of the events in sample space. Fig. 2 illustrates, for three such typical instances, the variation of the cumulative recognition score after every 100 samples for different values of λ . To restrict the size of the paper, the performance curves only for three instances are presented. Results obtained with self-supervised learning were compared with those for fully supervised and nonsupervised cases. It was revealed, under investigation, that the sequence corresponding to Fig. 2a provided a worse set of input events after the 5th step. The nonsupervised system, as expected, resulted in a poor performance where the large number of wrong classifications further weakened the already weak representative points. The reverse is the true for Fig. 2b where the sequence provided a better set of input samples after the 3rd step. For Fig. 2c, the first 100 patterns contained the best and resulted in higher initial recognition scores of 80% for fully supervised and 74% for nonsupervised learnings; the following sequences of worse events have reduced these scores to 75.5% and 67.75%, respectively.

As λ (in eqns. 8 and 9) increases/decreases, the dimension of the guard zones σ'_{nj} decreases/increases and the corresponding DPS-values are increased/decreased. Therefore the chance of correct/wrong samples correcting/vitiating the representative vectors is decreased/increased and the system performance accordingly approaches the nonadaptive/nonsupervised cases. Based on the mean-square distance, defining the mean-square error (Table 1) at every instance of the self-supervised curves, the curve corresponding to $\lambda = 2$ shows the best match with that of the fully supervised case. The entries shown in Table 1 are the average values computed over ten different sequences of inputs. The classifier corresponding to $\lambda = 2$ has been found, on average, to make available the highest proportion of correct to incorrect samples, so that after the several utterances have been dealt

Table 1: Mean-square distance (mean-square error) of self-supervised curves from fully supervised curve (averaged over ten observations)

λ^2	0.5	1	2	4	6	8
Mean-square distance	1.341	2.672	2.317	0.929	1.739	1.676

with by the classifier, the class-representative parameters are likely to approach their respective true values (as determined by the correct mean and standard deviations of the classes). The above results conform to our earlier findings [7], where the classifier was based on the 'minimum weighted distance function'. Tables 2a to c illustrate the confusion matrices of vowel recognition corresponding to Fig. 2a-c for both the cases of fully supervised (upper score) and optimum self-supervised (lower score) learnings. A figure in a cell denotes the number of times the machine took the same decision in recognising the vowel sounds for the two respective cases. The diagonal elements represent the number of utterances correctly identified. The confusion in vowel recognition, as expected from previous research [7-11], is seen to be restricted to two neighbouring classes.

8 Conclusions

The model of self-supervised learning algorithm with a classifier based on the fuzzy properties of patterns has been implemented to the real problem of vowel recognition.

System performance for different guard zones selected around the initial representative vectors is studied for different orders of input samples. With the shrinkage of the zone

boundaries, the system behaves like a nonadaptive recognition system, whereas the nonsupervised performance is approached by relaxing the boundaries. Optimum results are obtained when the semiaxes of the guard zones defined for the classes correspond to one half the standard deviations along the respective co-ordinate axes.

The self-supervision algorithm does not involve any concept of the theory of fuzzy sets, and hence can be used with any other classifier to improve the system's performance in recognising patterns. In deciding the relative merit between fuzzy and statistical classifiers, one should study the computational procedures involved and their respective efficacies in a given circumstance. For example, if we consider a Bayesian maximum likelihood classifier [1]; it uses more prior information about the system description and needs a large number of training samples in its design stage in order to evaluate the mean vectors, and magnitude and inverse of dispersion matrices corresponding to each of the classes.

Its decision-making stage involves a product (corresponding to each of the classes) of three matrices having dimensions $(1 \times N)$, $(N \times N)$ and $(N \times 1)$, respectively in order to classify an unknown pattern. The recognition scores for vowel and plosive sounds as obtained by such a nonadaptive statistical classifier using more prior information were found to be slightly better (2-3%) [5, 8, 10, 12] compared with those obtained by a fuzzy classifier reported in this paper. This fuzzy classifier, on the other hand, has lesser computational complexity and memory requirement for storing the representative parameters. It also gives satisfactory performance, even when the number of training samples is small enough to design a statistical classifier [8, 10]. Again, the fuzzy approach has been found to have more computational flexibilities because of the various operators and connectives, which can be exploited according to the extant problems [13].

9 Acknowledgments

The author wishes to thank Professors D. Dutta Majumder and A.K. Datta for their valuable discussion, and N.R. Ganguli and B. Mookherjee for processing the spectrograms. Thanks are also due to R.A. King for his interest in this work.

10 References

- MENDEL, J.M., and FU, K.S. (eds.): 'Adaptive, learning and pattern recognition system - theory and applications' (Academic Press, NY, 1970)
- YOUNG, T.Y., and CALVERT, T.W.: 'Classification, estimation and pattern recognition' (American Elsevier Publishing Co., London, 1974)
- ZADEH, L.A.: 'Outline of a new approach to the analysis of complex systems and decision processes', *IEEE Trans.*, 1973, SMC-3, pp. 28-44
- ALLEN, A.D.: 'Measuring the empirical properties of sets', *ibid.*, 1974, SMC-4, pp. 66-73
- PAL, S.K., and DUTTA MAJUMDER, D.: 'On automatic plosive identification using fuzziness in property sets', *ibid.*, 1978, SMC-8, pp. 302-308
- ZADEH, L.A., FU, K.S., TANAKA, K., and SHIMURA, M. (eds.): 'Fuzzy sets and their applications to cognitive and decision processes' (Academic Press, NY, 1975)
- PAL, S.K., DATTA, A.K., and DUTTA MAJUMDER, D.: 'A self-supervised vowel recognition system', *Pattern Recognition*, 1980, 12, pp. 27-34
- PAL, S.K.: 'Studies on the application of fuzzy set theoretic approaches in some problems of pattern recognition and man-machine communication by voice'. Ph.D. thesis, Calcutta University, 1978
- PAL, S.K., DATTA, A.K., and DUTTA MAJUMDER, D.: 'Adaptive learning in classification of fuzzy patterns - an application to vowels in CNC context', *Int. J. Syst. Sci.*, 1978, 9, pp. 887-897
- DUTTA MAJUMDER, D., DATTA, A.K., and PAL, S.K.: 'Computer recognition of Telugu vowel sounds', *J. Comput. Soc. India*, 1976, 7, pp. 14-20

Table 2: Confusion matrices for vowel recognition

		Actual class					
		I	E	δ	a:	O	U
Recognised as	I	106 117	3				
	E	64 51	163 151	10 12		6 7	15
	δ	2 4	30 39	51 49	26 27	1 1	2 2
	a:			8 8	57 58	5 5	
	O		9 9	2 2	6 4	133 136	23
	U		5 5	1 1		35 31	134 126

(i)

		Actual class					
		I	E	δ	a:	O	U
Recognised as	I	107 120	1				
	E	63 49	159 151	8 12		5 8	
	δ	2 3	33 40	53 49	30 27	2 1	2 2
	a:			7 8	52 56	5 5	
	O		11 10	3 2	7 6	132 135	15 23
	U		4 5	1 1		36 31	134 126

(ii)

		Actual class					
		I	E	δ	a:	O	U
Recognised as	I	99 117	3				
	E	72 49	186 142	11 11		11 7	1
	δ	1 6	14 48	51 50	28 29	6 1	2 2
	a:			7 8	55 57	2 5	
	O		4 9	2 2	6 3	129 132	11 11
	U		3 5	1 1		32 35	137 138

(iii)

(i) Corresponding to Fig. 2a, (ii) Corresponding to Fig. 2b, (iii) Corresponding to Fig. 2c
Upper score: fully supervised; lower score: optimum self-supervised ($\lambda = 2$)

- 11 PAL, S.K., and DUTTA MAJUMDER, D.: 'Effect of fuzzification on automatic vowel sound recognition'. Proceedings of fourth international joint conference on pattern recognition, Kyoto, Japan, 1978, pp. 1044-1046
- 12 DATTA, A.K., GANGULI, N.R., and RAY, S.: 'Recognition of unaspirated plosives - a statistical approach', *IEEE Trans.*, 1980, ASSP-28, pp. 85-91
- 13 JAIN, R.: 'Comments on 'Fuzzy set theory versus Bayesian statistics'', *ibid.*, 1978, SMC-8, pp. 332-333



Dr. Pal was born in Calcutta in 1950. He obtained B.Sc. (Hons.) in physics and B.Tech., M.Tech., and Ph.D. in radio-physics and electronics in 1969, 1972, 1974 and 1979, respectively, from the University of Calcutta. At present he is working in the Electrical Engineering Department, Imperial College, London, as a Commonwealth scholar '79 and is on leave from the post of computer engineer in the Indian Statistical Institute, Calcutta. His research interests include pattern recognition and image processing using fuzzy logic. He has published 35 papers and is a co-author of three edited books.

He is one of the reviewers of *Mathematical Reviews* (USA) in the fields of fuzzy sets, logic and applications, and a member of the Computer Society of India, the British Pattern Recognition Association, and of the IEEE.

Book review

Processing of visible language 2

P.A. Kolers, M.E. Wrolstad and H. Bouma
Nato Conference Series, Series II: Human factors
 Plenum Press, 1980, 616 pp., \$49.50
 ISBN: 0-306-40576-8

This volume is the second in a series of conference proceedings intended to bring together graphic designers, engineers and psychologists. In fact, the 37 papers range from the nature of writing systems, including Egyptian hieroglyphics (Meltzer), through computer conferencing (Baer and Turoff) to philosophical aspects of representation (Howard). Faced with the difficult task of knitting such diverse papers together, the editors have grouped them in sections with alternating emphasis on textual material and pictorial material. Like so many interdisciplinary collections, overall coherence suffers because of wide variation in the target audiences of the papers. Several tutorial reviews lack 'a message' and are rather too basic in their coverage. Many of the specialised papers are aimed at restricted audiences, and make few concessions to interdisciplinary communication.

In spite of these problems, the open-minded interdisciplinary reader could take away some important object lessons from this collection, both on specific points and on the different perspectives adopted by the various disciplines on common issues. The main area of overlap between the disciplines involves the human factors of the use of written or symbolic material where user understanding is critical. In the context of statistical graphs, Wainer discusses the incorporation of unnecessary graphic embellishments - 'chartjunk' - which detract from the simple and clear abstraction of underlying statistical relationships. Interestingly, Wainer, a psychologist, advocates the practical resolution of interposing graphical specialists between the originator of the material and the actual artist.

As an industrial designer, Doblin incorporates the same thing in the guise of 'decodability' within the much broader context of 'getting the message across', including the design context, power and credibility. The central feature of Doblin's analysis is his categorisation of messages and his model for

analysing the effectiveness of designs; this involves the questioning of various levels of the design and their consequences for the user.

A broad questioning process also underlies Wright's analysis of the usability of written material and textual literacy. Wright is a psychologist and her questions are derived from empirical research. Like Doblin, Wright advocates critical evaluation of draft material but, like Frase *et al.*, she emphasises the importance of direct empirical evaluation by systematic pretesting of documents on an appropriate sample of users.

A theme which emerges from Wright's analysis is that understanding depends not only on the structure and content of material, but also on the purpose for which someone is reading or using it. This important point recurs, in Perkins analysis of pictorial material in the form of matching presentation to the viewers, 'habits of information pickup', and Shebilske's, who stresses 'the active contribution of the reader'.

The volume emphasises technological developments in computerised presentation of text and graphics. The interests of the technologists and psychologists reflect the current drive towards more usable interfaces. Here, there are marked contrasts. Baecker, a computer scientist, describes the design of human-computer dialogue as an 'art form' in which advantageous system features assume prominence. Most psychologists prefer a more empirical approach. Sadly, some of the representatives of the latter, although technically sound, should be interpreted cautiously. As Brown *et al.* point out in one paper, particular dependent measures can drastically influence the conclusions of experiments on character recognition. Another paper by Treurniet is a case in point. This reports a study of character spacing on a VDU using a simple character search task. It is by no means clear that these results can be translated into firm generalisable guidelines.

In sum then, this is a disappointing collection, but it is punctuated with highlights. The interdisciplinary exchange is vital and such collections can *only* be expected to yield dividends if these disciplines try harder to understand each others' point of view.

P. BARNARD