

Classification using ML Methods and Neural Networks

by Edward Patch

Submission date: 29-Mar-2023 11:22PM (UTC+0100)

Submission ID: 200947905

File name:

11299_Edward_Patch_Classification_using_ML_Methods_and_Neural_Networks_495066_33532814.pdf
(728.71K)

Word count: 4050

Character count: 22127

Classification using ML methods and Neural Networks

1st Given Edward Patch

Software Engineering and Artificial Intelligence (of MSc Year 4)

Application of Machine Learning

University of Wales Trinity St. Davids (of Dr. Gordan Dickers)

Swansea, Wales

Student ID: 1801492

Abstract—Image classification is crucial in computer vision, impacting various fields like security, healthcare, and entertainment. This work explores and compares the performance of Support Vector Machine and Artificial Neural Network algorithms on the MNIST Fashion and CIFAR-10 datasets. Dimensionality reduction techniques, such as Linear Discriminant Analysis and Principal Component Analysis, were applied for preprocessing. The SVM with the RBF kernel achieved 84% accuracy on MNIST Fashion and 52% on C-10, whilst the ANN with a two-layer architecture reached 84.27% and 51.73% accuracy, respectively. Performance is heavily influenced by hyperparameter selection and dimensionality reduction techniques. Although results are specific to the datasets, the insights can potentially generalize to other image classification problems, informing the development of advanced algorithms and improving state of the art in computer vision.

Index Terms—SVM, ANN, Neural Networks, Artificial Intelligence, Deep Learning.

I. INTRODUCTION

A fundamental problem often presented in Artificial Intelligence (AI) is Image Classification in different subject areas. Many companies have leveraged Image Classification technology to enhance the functionality of their products. This paper explores the application of both traditional Machine Learning (ML) algorithms and the specialised subset of ML, Deep Learning (DL), which utilises Neural Networks (NN) to learn complex representations of high-dimensional data.

The study uses two popular datasets, Modified National Institute of Standards and Technology (MNIST) Fashion, obtained from Kaggle [1] and Canadian Institute for Advanced Research (CIFAR-10) constructed by Alex Krizhevsky [2], are available for the upcoming experiments to support the research.

Next, the MNIST Fashion dataset will aid with the fine-tuning and setup of the chosen ML classifier, Support Vector Machines (SVMs). Compared to Random Forest and Multinomial Logistic Regression (MLR), SVM is more capable of handling high dimensional space than the other two models. The interpretability and less prone to overfitting appear less prominent than in the alternative ML image classifications.

Once the first dataset is tuned and tested, the C-10 dataset is run through and compared with the MNIST dataset. However, this is not the full extent of how the experiment will take effect.

A detailed evaluation of MNIST Fashion and C-10 dataset Image Classification results are compared. The evaluation would use classification reports for identifying confusion matrices and Training and Validation that help identify any potential issues with the performance and average of the classifiers. The structure of the experiment will have an ML classifier and DL algorithm, ANN, to filter both MNIST and C-10 image classification results. When fine-tuning the ML classifier and DL algorithms, a comparison of the SVM test and trained data is cross-referenced with ANN. Thus, displaying any mistakes when handling the data helps identify any outliers within the data. This process also helps identify any potential amendments when the C-10 dataset is processed.

II. RELATED WORK IN LITERATURE

Preparing the dataset for the MNIST dataset requires advanced techniques to save time and gain better Image Classification results. A selection of related research papers aids the overall understanding, outlying any underlying problems or tips to benefit any future conduction of experiments.

The SVM's Support Vector Classifier (SVC) offers two fundamental parameters to SVM's allow non-linearity features to work best by constraining the 'C' parameter. Hence, it is more than 0 to adjust the penalty count, and the kernel influenced the overall accuracy of the test data. [3]

From the 'Feature Extraction Based on Deep Learning for Some Traditional Machine Learning Methods' by Aykut Cayir [4], the report demonstrates some findings of the SVM results for both MNIST Fashion and C-10 datasets. Results found in 'Identifying Critical Neurons in ANN Architectures using Mixed Integer Programming' by Mostafa ElAraby for MNIST Fashion achieves 97.8%, and C-10 achieves 39.13%. Mostafa ElAraby [5] experiments with the MNIST Fashion and CIFAR-10 dataset with the model using FC-3 and FC-4 connected layers to achieve 89.5% on Fashion MNIST and 72.3% C-10 datasets. Unfortunately, the author failed to provide the hyperparameters for reference.

Dimension and Feature Reduction (DaFR) plays a role in ML algorithms by reducing the dimension space of a large dataset. A significant problem, depending on the type of ML classifier and the hyper-parameter tuning, is that fitting the data may take too long and generate inaccurate results. DaFR has

two main categories: feature selection and feature extraction. Feature extraction creates a new set of features based on the original input and decreases the high-dimensional feature vector. Linear and non-linear are methods for feature extraction, where Principal Component Analysis (PCA) (Unsupervised) and Local Discriminant Analysis (LDA) (Supervised), where Auto-Encoder (AE) (Unsupervised) is a widely used non-linear method. [6] [7]

III. METHODOLOGY

The pipeline to construct the Image Classification experiment involves the preparation of the data, fine-tuning and final processing of the data. Using this methodology allows reasonable control over the test and training data, thus, giving additional hyper-parameters to fine-tune and analyse the data further.

After extensive research on DaFR, LDA appears to suit the MNIST Fashion dataset. Although PCA performs better at reducing dimensions, after some testing, LDA performs well at keeping essential data. LDA also considers the current features to give the algorithm a better understanding of necessary features. LDA and PCA similarly impacted the MNIST Fashion dataset's accuracy and performance. However, the PCA required more fine-tuning to increase accuracy results. Using the LDA prevented further testing of the SVC involving 'linear' and 'sigmoid' kernels. The reason to use DaFR is that the process exceeded forty-five minutes and may have taken too long to complete for the desired results. Experimentation of LDA found that the selection of nine classes was enough to reduce the dimensions without affecting the classes and impacting performance and accuracy in a good way.

The SVM SVC image classifier experiment uses the following hyperparameters:-

```
Kernel='RBF'
C=1
Gamma=0.001
Random_State=85
Decision_Function_Shape='OVR'
```

The selection of the hyperparameters is the most optimised on the MNIST Fashion dataset. The Kernel parameter focuses on how the algorithm transforms input data to a higher-dimensional representation.

The RBF kernel is the best combination of hyperparameters for the MNIST Fashion dataset. The parameters involve the C parameter is set to 1 preventing any misclassified data points, γ is set to 0.001 to control the scaling of RBF, random_state seed is set to 85, and decision_function_shape performs better with One versus Rest (OvR) than One versus One (OvO), achieving an accuracy score of eighty-four per cent. However, there was an accuracy score of eighty-five per cent, with C being the value 10. Although it appeared to have a higher accuracy score, it had lower precision, recall and F1-Scores than other configurations.

ANN usually consists of three layers: input, hidden and output. A ANN input layer model was setup using Tensorflow's Kera's version of the library. Using Kera's Sequential function

call, two layers worked best with two layers: input and output. It is important to establish how the hyperparameters work. The 'units' parameter controls the number of neurons used to process the hyperparameters. The 'use_bias' parameter is a boolean datatype with the purpose of influencing the weighted sums of inputs within a neuron. The activation parameter is how the neurons should function. This activation function introduces non-linearity into the model by setting all negative input values to zero. 'kernel_initializer' is a parameter that controls the weights of neurons within the model layers.

The layers are as follows:

```
tf.keras.layers.Dense(units=64,
    use_bias=False, activation='relu',
    input_shape=(9,)),
    kernel_initializer=tf.keras.
    initializers.
    RandomNormal(mean=0.01, stddev=0.1)),
    tf.keras.layers.Dense(units=25,
    input_shape=(9,)), use_bias=True,
    activation='sigmoid'),
    tf.keras.layers.Dense(10,
    activation='softmax',
    kernel_regularizer=tf.keras.
    regularizers.l2(0.01))
```

The input layer uses sixty-four units, false bias, 'relu' [8] activation, nine features for the input shape and a kernel initialiser; focusing on the keras's RandomNormal initialiser with mean set as 0.01 and standard deviation set as 0.1. The hidden layer uses twenty-four units, true bias, 'sigmoid' activation and nine input layers. An issue that may happen with Sigmoid activation function is susceptible to the 'Vanishing Gradient Problem' [9]. The output layer uses ten units, 'softmax' [10] activation, and L2 regulariser set as 0.01. The L2 regularisers has a loss of 0.5%. This achieves an accuracy score of 84.25% and displays a learning trend on the '(ANN) Training and Validation Loss for MNIST Fashion Dataset' 9 shows a steep incline suggesting a significant fall of data loss. Overfitting and underfitting is reduced and Figure 9 shows the Training Loss line moved closer to the Validation Loss line. However, after the first epoch of the training and validation, it steadies and follows the validation loss. This result shows that the data points are fitting relatively closely until the fifth epoch. The trend is then overlapping the Validation loss indicating that the algorithm is performing well.

The initial drop in data loss could be attributed to the algorithm's initial efforts to understand and adapt to the data points. Overall, the NN model with the specified architecture and hyperparameters performs well on the MNIST Fashion dataset, achieving a reasonable accuracy score while reducing overfitting and underfitting.

IV. RESULTS - DIMENSION AND FEATURES REDUCTION

Figures (1, 2) show the MNIST Fashion and C10 datasets projected onto the first two principal components identified by PCA. The scatter plot visualises the ten available classes in the reduced-dimensional space.




Fig. 1. PCA of MNIST Fashion Dataset




Fig. 3. LDA of MNIST Fashion Dataset




Fig. 2. PCA of C-10 Dataset




Fig. 4. LDA of C-10 Dataset

Figures (3, 4) present a scatter plot of the MNIST Fashion dataset projected onto the first two linear discriminants identified by LDA. The field displays the ten available classes in the reduced-dimensional space, visualising the dataset's separability concerning these classes.

V. RESULTS - CONFUSION MATRIX HEATMAPS

Figures (5, 6) contain a SVM Confusion Matrix Heatmap for the MNIST Fashion and C-10 datasets detailing the overall fitting of the Validation Training.

Figures (7, 8) illustrate a ANN Confusion Matrix Heatmap for MNIST Fashion and C-10 datasets.

VI. RESULTS: VALIDATION AND TRAINING LOSS

Figures (9, 10) demonstrates the learning curvature of each between the training and validation data points by using the loss for reference for both datasets.

VIII. DISCUSSION

Referring to the DaFR, LDA was practical on the MNIST Fashion dataset. A comparison of figures (3, 4) reveals that the MNIST Fashion dataset has distinct and separable features, enabling LDA to filter and reduce ambiguous features efficiently. On the other hand, the C-10 dataset contains more complex and diverse images, possibly due to the similarity between classes or the more significant number of pixels. While PCA, an unsupervised method, visualises and compresses data without pairing up given features, it may not necessarily outperform LDA on the C-10 dataset. PCA could be useful for preprocessing data for the SVC and ANN setup, depending on the underlying problem.

Tables (I, II): Comparison of the Classification Reports of both MNIST Fashion and C-10 dataset with SVC indicates a few factors of interest. These factors include each classifi-




Fig. 5. SVM Confusion Matrix Heatmap for MNIST Fashion Dataset - Inspired by 'Heart Failure Prediction: SVM and ANN' [11]




Fig. 6. SVM Confusion Matrix Heatmap for C10 Dataset - Inspired by 'Heart Failure Prediction: SVM and ANN' [11]

cation's precision, recall, F1-Score and support. Furthermore, Classification Reports have the accuracy, macro average and weighted average to indicate the overall SVC performance. After close observation, the MNIST Fashion and C-10 datasets have ten classes. However, they differ in class types. The MNIST dataset holds 'T-Shirt/Top', 'Trouser', 'Pullover', 'Dress', 'Coat', 'Sandal', 'Shirt', 'Sneaker', 'Bag' and 'Ankle Boot'. In contrast, the C-10 dataset contains a mixture of classifications, including 'Airplane', 'Automobile', 'Bird', 'Cat', 'Deer', 'Dog', 'Frog', 'Horse', 'Ship' and 'Truck'.

The categories within the Classification Report are defined as the following:

- 'Precision' is the number of true positives divided by the total number of true and false positives. It measures how many of the predicted positive samples are positive.




Fig. 7. ANN Confusion Matrix Heatmap for MNIST Dataset - Inspired by 'Heart Failure Prediction: SVM and ANN' [11]




Fig. 8. ANN Confusion Matrix Heatmap for C-10 Dataset - Inspired by 'Heart Failure Prediction: SVM and ANN' [11]

- 'Recall' is the number of true positives divided by the total number of true positives and false negatives. It measures how many positive samples are correctly predicted as positive.
- 'F1-Score' is the harmonic mean of precision and recall. It provides a balanced measure of precision and recall.
- 'Support' is the number of samples in each class.

The Classification Reports provide accuracy, macro-average, and weighted average values to measure the overall performance of the classifiers. Precision, recall, and F1-score were used to evaluate the performance of the MNIST Fashion and C-10 datasets. MNIST Fashion had a more comprehensive range of precision values (65-99%) compared to C-10 (37-61%), indicating a more apparent range of images. Both




Fig. 9. ANN Training and Validation Loss for MNIST Fashion Dataset - Inspired by 'Heart Failure Prediction: SVM and ANN' [11]




Fig. 10. ANN Training and Validation Loss for C-10 Dataset - Inspired by 'Heart Failure Prediction: SVM and ANN' [11]

datasets performed well regarding the recall, suggesting little overfitting or underfitting. The F1-score for MNIST Fashion ranged from 60% to 97%, while for C-10, it went from 40% to 62%. The Ship classification had relatively good accuracy/recall compared to the others, based on the mean of precision and recall. The support for both datasets is fairly similar, with MNIST Fashion having 36,000 samples and C-10 having 30,000 samples. The overall accuracy is 84% (MNIST Fashion) and 52% (C-10).

Figures (5, 6, 7, 8) provide the Confusion Matrix Heatmaps for SVM's classifier results. The Confusion Matrix demonstrates where the classification may have gone wrong, giving more information on what techniques may benefit from further tweaking and training. The heatmap shows the depth of

VII. RESULTS - CLASSIFICATION REPORT

Tables (I, II) show the performance of an SVM's SVC image classifier on the MNIST Fashion dataset in a Classification Report format. The table includes class-specific precision, recall, F1-score, support, and overall accuracy, providing a comprehensive summary of the model's performance of test data training.

TABLE I
MNIST FASHION - CLASSIFICATION REPORT FOR THE 'SVM.SVC' MODEL.

Class	Precision	Recall	F1-Score	Support
T-Shirt/Top	0.78	0.83	0.80	3565
Trouser	0.99	0.95	0.97	3602
Pullover	0.78	0.70	0.74	3622
Dress	0.81	0.88	0.85	3578
Coat	0.72	0.80	0.75	3580
Sandal	0.91	0.90	0.90	3635
Shirt	0.65	0.56	0.60	3578
Sneaker	0.88	0.91	0.89	3655
Bag	0.94	0.95	0.94	3622
Ankle Boot	0.93	0.92	0.92	3563
Accuracy			0.84	36000
Macro Avg.	0.84	0.84	0.84	36000
Weighted Avg.	0.84	0.84	0.84	36000

Accuracy Score: 84%

TABLE II
CIFAR-10 - CLASSIFICATION REPORT FOR THE SVM.SVC MODEL.

Class	Precision	Recall	F1-Score	Support
Airplane	0.57	0.57	0.57	3025
Automobile	0.61	0.57	0.59	2997
Bird	0.44	0.40	0.42	3018
Cat	0.37	0.44	0.40	2920
Deer	0.47	0.45	0.46	2980
Dog	0.46	0.45	0.45	3003
Frog	0.55	0.55	0.55	3034
Horse	0.60	0.51	0.55	3044
Ship	0.59	0.64	0.62	3002
Truck	0.57	0.60	0.58	2977
Accuracy			0.52	30000
Macro Avg.	0.52	0.52	0.52	30000
Weighted Avg.	0.52	0.52	0.52	30000

Accuracy Score: 52%

support a class may have had in both True and Predicted labels. For our guide to describe true positives, high (SVM: ≥ 2500 — ANN: ≥ 400), medium (SVM: ≥ 1500 and < 2500 — ANN: ≥ 300 and < 400), and low (< 300), and for the false positives, high (SVM: > 500 — ANN: > 50) and low (SVM: ≤ 500 — ANN: ≤ 50). MNIST Fashion has nine high true positives and one medium true positive. According to the figure 5, the classifier does well at identifying distinctive images. Still, it fails with high false positives on 'Pullover', mistaking it highly for 'Coat' and 'Shirt' for the low side. This occurrence happens similarly for 'T-Shirt/Top', mistaking it for 'Shirt'. Two of MNIST Fashion's top performers are 'Trouser' and 'Bag', with high true and low false positives. ANN has similar results with fewer samples. Seven high true positives, two medium true positives and one low true positive are present. The 'trouser' classifier has the highest true and very low false positives.

Figure 6 has six medium true positives and four low true positives. After close examination, one high false positive appears, regardless of the low number of true positives. This

analysis shows that the classifications are relatively spread, although the type has more good than poor decisions. For example, the classifier is successful at selecting distinctive labels, which means the spread of choice could be impacted by the colour of the dataset. Greyscale tends to eliminate any lighting issues and adds crispier outlines. Camera angles may also be at fault, with 'automobile' and 'truck'. The 'ship' classification gets mistaken for an 'airplane'. Perhaps the classifier is identifying the image's background, seeing the sky's reflection, rather than judging the outline. The ANN fails to have medium or high true positives. However, there is no prominent overfitting/underfitting happening. The results indicate common misclassifications previously mentioned, such as 'airplane' and 'ship' and 'truck' and 'automobile' getting mixed up, backing up some of the points previously mentioned.

ANN performs and achieves higher accuracy, with MNIST Fashion being 84% and C-10 being 51.73%. C-10 has a slightly higher accuracy than SVM.SVC. Figures (9, 10) has a few performance differences. The MNIST Fashion will have an initial loss drop and then starts to tighten up against the validation loss by the fourth epoch. The C-10 dataset, on the other hand, would drop and bypass the validation loss by the second epoch. After that, the NN struggles to fit and continues to underfit til the tenth epoch.

These tests are different to MNIST Fashion and C-10 datasets because the C-10 dataset is significantly more extensive and has another resolution, a diverse range of classifications, and RGB rather than Greyscale. The hyperparameters may have been adjusted to suit MNIST Fashion rather than C-10 dataset, so without further investigation, it is unclear whether the algorithms used could cut it or not. For future experiments, experiment with PCA instead of LDA as PCA handles overlapping much better, and different kernels are available for SVM.SVC and change the C parameter to be deeper than 1. The ANN may benefit significantly from more layers or different hyperparameters and activation parameters. This suggestion might allow C-10 to generate better performance and accuracy.

IX. CONCLUSION

In conclusion, this project aimed to explore and compare the performance of different image classification algorithms on the MNIST Fashion and C-10 datasets. The two algorithms used were SVM and ANN. The SVM algorithm used the RBF kernel and achieved an accuracy score of 84% on the MNIST Fashion dataset and 52% on the C-10 dataset. The ANN used a two-layer architecture with ReLU and Sigmoid activation functions, achieving an accuracy score of 84% on the MNIST Fashion dataset and 51.73% on the C-10 dataset.

It was found that the LDA dimensionality reduction technique worked well with the MNIST Fashion dataset, while the PCA technique required more fine-tuning to increase accuracy results. Additionally, the RBF kernel performed the best out of the four SVM kernels tested. For the ANN, using L2 regularization helped reduce overfitting and underfitting.

The C-10 dataset offers a different range of images with colour information. This extra information could make it harder for the current SVM.SVC hyperparameters to identify the images. For instance, a dog could look like a cat dependent on the camera angle, colours and image quality. The MNIST Fashion also has similar images in most pairs, and greyscale can help dim images and make the shape more definitive. While MNIST Fashion has a more comprehensive range of precision values than C-10, it is essential to note that the actual precision values for both datasets are relatively high. This observation indicates that the models are performing well on both datasets overall, despite differences in the consistency of their performance across different classes.

Overall, this project provides insights into the performance of different algorithms on image classification tasks and highlights the importance of fine-tuning hyperparameters to achieve optimal results. These insights can be used to guide future research in image classification and aid in the development of more advanced algorithms. The results in this experiment for SVM achieve 84.25% for MNIST Fashion and 52%. Whereas, A. Cayir [4] has a 97.8% for MNIST Fashion and 39.13% for the C-10 dataset. This ANN experiment achieves 84.25% and 51.73%, whereas M. ElAraby [5] achieves 89.5% on Fashion MNIST and 72.3% C-10 datasets. Reviewing the comparison of these results shows that even though A. Cayir [4] achieves a higher impact for the SVM classifier, though the CIFAR-10 suffers massively. The experiment conducted by M. ElAraby [5] achieves higher results in both datasets. However, it does not state the performance of the models, so overfitting/underfitting could be happening. Future work could involve fine-tuning and changing layers of ANN, with the C-10 dataset being the primary tuner and the MNIST Fashion dataset for comparison.

X. TERMINOLOGY

List of terminologies used in this document:-

- ML - Machine Learning.
- DL - Deep Learning.
- NN - Neural Networks.
- AI - Artificial Intelligence.
- MNIST - Modified National Institute of Standards and Technology Database.
- CIFAR-10 - Canadian Institute for Advanced Research.
- SVM - Support Vector Machine.
- MLR - Multinomial Logistic Regression.
- ANN - Artificial Neural Network.
- DaFR - Dimension and Feature Reduction.
- PCA - Principal Component Analysis.
- LDA - Local Discriminant Analysis.
- RBF - Radial Basis Function.

XI. REFERENCE LIST

- [1] Zalando Research, "Fashion MNIST." [Online]. Available: <https://www.kaggle.com/datasets/zalando-research/fashionmnist>

- [2] Alex Krizhevsky, Vinod Nair, and Geoffrey Hinton, "CIFAR-10 and CIFAR-100 datasets." [Online]. Available: <https://www.cs.toronto.edu/~kriz/cifar.html>
- [3] K. M. Buddhiraju and I. A. Rizvi, "Comparison of CBF, ANN and SVM classifiers for object based classification of high resolution satellite images," in *2010 IEEE International Geoscience and Remote Sensing Symposium*, Jul. 2010, pp. 40–43, iSSN: 2153-7003.
- [4] A. Çayır, I. Yenidoğan, and H. Dağ, "Feature Extraction Based on Deep Learning for Some Traditional Machine Learning Methods," in *2018 3rd International Conference on Computer Science and Engineering (UBMK)*, Sep. 2018, pp. 494–497.
- [5] M. ElAraby, "Identifying Critical Neurons in ANN Architectures using Mixed Integer Programming," *DeepAI*, Feb. 2020. [Online]. Available: <https://deepai.org/publication/identifying-critical-neurons-in-ann-architectures-using-mixed-integer-programming>
- [6] H. AlSaeed, N. Hewahi, and R. Ksantini, "Dimension Reduction Techniques for Image Classification," in *2022 International Conference on Innovation and Intelligence for Informatics, Computing, and Technologies (3ICT)*, Nov. 2022, pp. 358–365, iSSN: 2770-7466.
- [7] M. Asadur Rahman, M. Faisal Hossain, M. Hossain, and R. Ahmmed, "Employing PCA and t-statistical approach for feature extraction and classification of emotion from multichannel EEG signal," *Egyptian Informatics Journal*, vol. 21, no. 1, pp. 23–35, Mar. 2020. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S1110866519301720>
- [8] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet Classification with Deep Convolutional Neural Networks," in *Advances in Neural Information Processing Systems*, F. Pereira, C. J. Burges, L. Bottou, and K. Q. Weinberger, Eds., vol. 25. Curran Associates, Inc., 2012, p. 3. [Online]. Available: <https://proceedings.neurips.cc/paper/2012/file/c399862d3b9d6b76c8436e924a68c45b-Paper.pdf>
- [9] Z. Hu, J. Zhang, and Y. Ge, "Handling Vanishing Gradient Problem Using Artificial Derivative," *IEEE Access*, vol. 9, pp. 22 371–22 377, 2021.
- [10] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*. MIT Press, 2016.
- [11] "Heart Failure Prediction: SVM and ANN." [Online]. Available: <https://kaggle.com/code/surajjha101/heart-failure-prediction-svm-and-ann>

Classification using ML Methods and Neural Networks

GRADEMARK REPORT

FINAL GRADE

78/100

GENERAL COMMENTS

Instructor

A good abstract with summary results. Well done.

Your related work/ lit review identifies useful related work. Good. References are well used and of a good quality, well done again.

The experiments include effects and advantages of dimensionality reduction. Nice to see this.

Results are generally in keeping with the literature and your discussions of these is good. To help the reader further, try to include bar charts as well as tables to summarize the results and easily aid comparisons. For the ANN please also provide the same metrics as tables I and II.

Nice to see you grouping into ranges to aid discussion and comparisons.

The report is very well structured and generally well written.

Well done.

PAGE 1



Comment 1

SAYS WHO - REF THIS

PAGE 2

PAGE 3

GRADING FORM: SMU GENERIC_GJ

EDWARD PATCH

0

RESEARCH-INFORMED LITERATURE

KNOWLEDGE AND UNDERSTANDING

ANALYSIS

PRACTICAL APPLICATION AND DEP.

SKILLS FOR PROFESSIONAL PRAC.