



**A DEEP LEARNING APPROACH FOR ROBUST CROP CLASSIFICATION UNDER
MIXED-PIXEL CONDITIONS IN HYPERSPPECTRAL IMAGES
SIMPLIFIED RESEARCH MANUSCRIPT BASED ON THE SUPPLIED SYNTHETIC
EXPERIMENT**

Praveenkumar V¹

*Assistant Professor, Department of Information Technology, Sri Venkateswara College of
Engineering, Sriperumbudur, Tamilnadu, India,
praveen.k930@gmail.com*

N Gobalakrishnan²,

*Professor, Department of Computer Science and Engineering, Saveetha Engineering College,
Thandalam, Chennai, Tamilnadu, India
gobalakrishnanse@gmail.com*

Arunnachalam M³

*Assistant Professor, Department of Information Technology, Sri Venkateswara College of
Engineering, Sriperumbudur, Tamilnadu, India
arunmurugan2805@gmail.com*

Abstract

Hyperspectral imaging records reflectance across many narrow wavelength bands and can help distinguish crop types. However, a pixel may contain signals from more than one crop, soil, shadow, or other material. This paper reviews the supplied synthetic benchmark for a mixture-aware spectral-spatial fusion concept called MA-SSFN. The experiment is described as using 48 abstract spectral features and 16 Salinas-inspired class labels; it does not establish the use of official Salinas imagery. Five methods were compared using overall accuracy (OA), macro-F1, Cohen's kappa, and macro recall. RBF-SVM achieved the highest benchmark scores (OA = 0.939643; macro-F1 = 0.939509; κ = 0.935619). In the ablation study, spectral features combined with smoothed context achieved the highest scores (OA = 0.93643; macro-F1 = 0.93631). The full spectral-context-mixture configuration performed less well. These results do not show that the fusion-inspired model is superior. Because raw data, implementation details, repeated runs, and real-scene validation are not supplied, the findings should be treated as descriptive results from a synthetic experiment.

Keywords: hyperspectral imaging; crop classification; mixed pixels; spectral-spatial learning; synthetic data; support vector machine; ablation study

1. Introduction

Crop maps can support agricultural planning, field monitoring, and crop-condition assessment. The practical value of agricultural technologies also depends on whether farmers can adopt and use them sustainably (Uma Rani & Padmalosani, 2019). Hyperspectral sensors measure reflectance in many narrow bands, potentially revealing differences related to pigments, canopy structure, water, and biochemical composition. These data are high-



dimensional, and their values can vary with illumination, atmosphere, crop stage, and sensor characteristics (Landgrebe, 2002).

A mixed pixel occurs when a pixel contains more than one material—for example, crop vegetation together with soil, weeds, shadow, or a neighbouring crop. Spectral mixture analysis describes an observation as a combination of component spectra, although real mixtures can be more complex than a simple linear combination (Keshava & Mustard, 2002; Somers et al., 2011). This makes mixed-pixel classification an important test for crop-mapping methods.

Conventional classifiers remain essential comparators. Support vector machines can learn nonlinear decision boundaries, while tree ensembles capture nonlinear feature interactions (Cortes & Vapnik, 1995; Breiman, 2001). Deep-learning methods can learn useful representations from data, and spectral–spatial models aim to use both spectral signatures and local context (Chen et al., 2014; Roy et al., 2020). However, model complexity alone does not guarantee better results (Audebert et al., 2019; Li et al., 2019).

This study therefore evaluates the numerical evidence provided for an MA-SSFN-inspired Fusion MLP. It compares the supplied benchmark and ablation results, identifies what they support, and outlines the additional tests needed before claims about real agricultural use can be made.

2. Related Work

2.1 Hyperspectral data and mixed pixels

Hyperspectral images contain detailed spectral information but often include correlated bands, limited labels, and within-class variation. Dimensionality-reduction methods such as principal component analysis can summarize correlated variables, although high variance does not necessarily mean high class discrimination (Jolliffe & Cadima, 2016). Mixture analysis provides a framework for describing subpixel composition, but endmember variability and measurement conditions complicate the task (Keshava & Mustard, 2002; Somers et al., 2011).

2.2 Conventional and deep-learning classifiers

SVMs and tree ensembles are useful baselines because they provide alternative ways to model nonlinear class boundaries (Breiman, 2001; Camps-Valls & Bruzzone, 2005). Neural networks can learn hierarchical representations (LeCun et al., 2015; Goodfellow et al., 2016). Research on hyperspectral classification includes CNNs, hybrid spectral–spatial networks, autoencoders, and other architectures (Chen et al., 2014; Ma et al., 2016; Roy et al., 2020). Reviews emphasize that data quality, sample size, and evaluation design affect comparisons among methods (Paoletti et al., 2019; Zhu et al., 2017).

Other model families, including recurrent networks, gated recurrent units, and encoder–decoder systems, may be appropriate when the data structure supports them (Hochreiter & Schmidhuber, 1997; Cho et al., 2014; Hinton & Salakhutdinov, 2006). AI is also being explored for environmental monitoring in controlled agricultural systems such as hydroponics, illustrating a related application of data-driven methods to agricultural management (Senthil Pandi et al., 2025). These studies provide broader agricultural context;



they do not validate the crop-classification scores reported in the present synthetic experiment. Their appearance in conceptual diagrams does not mean that they were tested in the supplied benchmark.

3. Materials and Methods

3.1 Data and study scope

This paper is a secondary analysis of the tables and diagrams in the supplied document. The benchmark is described as a synthetic experiment with 48 abstract spectral features and 16 Salinas-inspired class labels. The available source does not establish that official Salinas image pixels or measured reflectance spectra were used. The scores are therefore reported as synthetic benchmark results, not as validation on the Salinas scene.

Important details are unavailable: the numerical feature matrix, class-wise sample counts, random seed, data split, model hyperparameters, mixture proportions, raw predictions, and confidence intervals. No new model was trained for this rewrite, and no unsupported performance values have been added.

3.2 Mixed-pixel concept

A common conceptual mixture model is $x = \sum a_k s_k + \epsilon$, where x is the observed spectrum, s_k is the spectrum of component k , a_k is its contribution, and ϵ represents noise or unexplained variation. In a simple linear model, contributions are nonnegative and sum to one (Keshava & Mustard, 2002). This equation explains the idea of synthetic mixing; it is not a claim about the exact code used to create the supplied data.

3.3 MA-SSFN-inspired model

The model concept combines three inputs: spectral features, smoothed local context, and a mixture-awareness descriptor. The branches are intended to be combined before a classifier predicts a class. The source reports a “MA-SSFN-inspired Fusion MLP,” but does not provide enough implementation detail to reproduce a complete spectral–spatial network. The model name is therefore retained as “fusion-inspired MLP” when discussing the results.

3.4 Evaluation metrics

The benchmark reports overall accuracy, macro-F1, Cohen’s kappa, and macro recall. Overall accuracy is the proportion of evaluated samples classified correctly. Macro-F1 averages class-level F1 scores, while macro recall averages recall across classes. Cohen’s kappa adjusts observed agreement for expected agreement based on class distributions (Cohen, 1960). These metrics should be considered together because each describes a different aspect of classification quality (Sokolova & Lapalme, 2009). Accuracy estimates also depend on representative reference labels and a defensible sampling design (Foody, 2002; Congalton & Green, 2019).

4. Results

Table 1 presents the supplied benchmark values. RBF-SVM has the highest reported value for all four metrics. Its overall accuracy is 0.939643 (93.9643%), compared with 0.932857 (93.2857%) for Spectral MLP and 0.919643 (91.9643%) for the MA-SSFN-inspired Fusion MLP. The RBF-SVM exceeds the Fusion MLP by 0.020000, or 2.00 percentage points, in overall accuracy. These differences are descriptive; the source does not provide repeated runs or uncertainty estimates.



Table 1. Benchmark performance of crop-classification methods

Method	Overall accuracy	Macro-F1	Cohen's kappa	Macro recall
RBF-SVM	0.939643	0.939509	0.935619	0.939643
Random Forest	0.888214	0.888949	0.880762	0.888214
Extra Trees	0.891071	0.891771	0.883810	0.891071
Spectral MLP	0.932857	0.932594	0.928381	0.932857
MA-SSFN-inspired Fusion MLP	0.919643	0.919960	0.914286	0.919643

Note. Values are proportions from the supplied synthetic experiment. Multiply by 100 to express them as percentages. They are not results from a verified real-scene Salinas evaluation.

Table 2 compares the three supplied feature configurations. Adding smoothed context to the spectral branch increases OA from 0.93286 to 0.93643 and macro-F1 from 0.93259 to 0.93631. The full spectral-context-mixture configuration scores lower: OA = 0.92143 and macro-F1 = 0.92116. The full configuration is 1.50 percentage points lower in OA than spectral plus smoothed context.

Table 2. Ablation results for feature configurations

Configuration	Overall accuracy	Macro-F1
Spectral branch only	0.93286	0.93259
Spectral + smoothed context	0.93643	0.93631
Full spectral-context-mixture fusion	0.92143	0.92116

Note. The spectral-plus-smoothed-context configuration has the highest supplied ablation scores. The mixture-awareness descriptor does not improve the reported point estimates in this synthetic experiment.

The results suggest that local context may help in this particular synthetic setup, while the additional mixture-awareness descriptor does not improve the reported scores. The source does not provide enough information to identify why. Possible explanations—such as redundant features, unsuitable scaling, or limited training data—remain hypotheses rather than established findings.

5. Interpretation of the Source Figures

5.1 Wavelength markers and plant indicators

Hyperspectral Reflectance and Biophysical Indicators Wavelength Markers and Their Interpretation




Figure 2 summarizes the wavelength markers shown in the source illustration. The figure groups markers across visible, red-edge, near-infrared (NIR), and shortwave-infrared (SWIR) regions and labels them in relation to pigments, stress, water, canopy structure, and plant compounds. These are candidate associations shown in a schematic figure, not measured evidence of crop condition or yield.

Table 3. Wavelength markers and interpretations shown in the source illustration


Indicator / feature	Wavelengths (nm)	Interpretation shown
Nitrogen-related indicator	405	Nitrogen-related indicator
Pigments	450, 515, 550, 570, 650	Pigment-related response
Chlorophyll absorption	687	Chlorophyll absorption feature
Light-use efficiency	490, 531	Light-use-efficiency indicators
Plant stress / senescence	705, 720; 700–740 region	Stress-related region
Biophysical / yield-related indicators	687, 760, 855, 1045, 1100	Biophysical or yield-related properties
Water	970, 1180, 1245, 1450, 1760, 1950, 2050, 2145, 2173	Water absorption regions
Lignin, cellulose, plant litter	1548, 1620, 1690, 2025, 2133, 2205	Structural-compound / litter-related regions
Heavy-metal stress (as labelled)	1650	Label shown in the source illustration

Plant stress (as labelled)	2295	Label shown in the source illustration
----------------------------	------	--

Note. These labels are transcribed from the supplied schematic. They should not be treated as uniquely diagnostic wavelengths without calibration and independent plant measurements.

The crop-growth-stage reflectance curves in the source are also schematic. Their axes show approximately 350–2350 nm and 0–60% reflectance, but the values in parentheses beside crop labels are undefined. Therefore, neither the wavelength illustration nor the crop curves can establish actual spectral differences, classification accuracy, or yield effects.

5.2 Deep learning and chemometrics diagrams



The source diagram compares deep-learning workflows with traditional machine-learning and chemometric workflows. It lists CNN and DAE among deep-learning preprocessing examples, CNN for feature extraction, and CNN, RNN, LSTM, and SAE for modelling. The traditional workflow lists smoothing and spectral preprocessing methods, feature-selection approaches such as PCA and CARS, and models including PLS, SVR, SVM, KNN, LDA, and decision trees. These are examples of possible methods from the present benchmark.

The architecture illustration also shows CNN, RNN, an LSTM-like gated cell, and an encoder–decoder. CNNs can learn local patterns, recurrent models represent ordered sequences, and encoder–decoder models learn compact representations and reconstruct or

predict outputs (Goodfellow et al., 2016; Hochreiter & Schmidhuber, 1997). The supplied benchmark does not report a direct test of these four architectures.

6. Research Map

Figure 1 gives a concise map of the study. It links the mixed-pixel problem to the supplied feature representation, comparison methods, reported metrics, leading benchmark result, ablation result, and key limitation. The map emphasizes that the strongest reported benchmark is RBF-SVM, while spectral features plus smoothed context lead the ablation.

Figure 1. Simple map of the research design and main findings




Figure 1. Simple map of the research design and main findings.

7. Discussion

The main finding is that the more complex fusion-inspired MLP did not outperform RBF-SVM in the supplied benchmark. This does not prove that SVMs are always better; it shows only that RBF-SVM led this comparison. The spectral-plus-smoothed-context configuration performed best in the ablation, suggesting that contextual information may be useful when applied carefully.

Context smoothing can reduce isolated noise, but it can also blur field boundaries and mix neighbouring classes. Future experiments should test performance separately at field edges and within fields (Pedregosa et al., 2011).

7.1 Limitations

The study is limited by missing raw features, sample counts, data-generation code, mixture settings, train/test split, hyperparameters, class-wise metrics, confusion matrices, repeated runs, and uncertainty intervals. The fusion-inspired model is not described in enough detail for exact replication. The spectral figures are illustrative, and the source does not establish official Salinas imagery as the data source. The results therefore support a descriptive comparison only.



A stronger follow-up study should publish the feature-generation code, class definitions, random seeds, mixture proportions, noise model, preprocessing steps, model configurations, and evaluation scripts. It should test a range of mixture severities, including crop–soil and crop–crop mixtures, and clearly define whether the target is the dominant crop or fractional composition.

Training, validation, and test sets should be separated to prevent spatial leakage. All methods should receive a documented tuning budget, and multiple random seeds should be used. Report means and uncertainty intervals alongside OA, macro-F1, macro recall, kappa, class-wise precision and recall, and confusion matrices. Performance should be grouped by class frequency and mixture severity. Finally, test the selected model on a public hyperspectral scene using its official image cube and ground-truth labels, followed where possible by independent-field or independent-date validation (Congalton & Green, 2019; Foody, 2002).

9. Conclusion

This simplified analysis preserves the supplied results while clarifying their limits. RBF-SVM achieved the strongest benchmark performance (OA = 93.9643%; macro-F1 = 93.9509%; κ = 93.5619%). In the ablation study, spectral features plus smoothed context achieved the highest scores (OA = 93.643%; macro-F1 = 93.631%). Adding the mixture-awareness descriptor did not improve the reported point estimates. These findings are based on a synthetic experiment and do not establish real-scene crop-classification robustness. The next step is a fully documented, repeated, spatially independent evaluation using measured hyperspectral data.

References

1. Audebert, N., Le Saux, B., & Lefèvre, S. (2019). Deep learning for classification of hyperspectral data: A comparative review. *IEEE Geoscience and Remote Sensing Magazine*, 7(2), 159–173. <https://doi.org/10.1109/MGRS.2019.2912563>
2. Breiman, L. (2001). Random forests. *Machine Learning*, 45, 5–32. <https://doi.org/10.1023/A:1010933404324>
3. Camps-Valls, G., & Bruzzone, L. (2005). Kernel-based methods for hyperspectral image classification. *IEEE Transactions on Geoscience and Remote Sensing*, 43(6), 1351–1362. <https://doi.org/10.1109/TGRS.2005.846154>
4. Chen, Y., Lin, Z., Zhao, X., Wang, G., & Gu, Y. (2014). Deep learning-based classification of hyperspectral data. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 7(6), 2094–2107. <https://doi.org/10.1109/JSTARS.2014.2329330>
5. Cho, K., van Merriënboer, B., Gulcehre, C., Bahdanau, D., Bougares, F., Schwenk, H., & Bengio, Y. (2014). Learning phrase representations using RNN encoder–decoder for statistical machine translation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing* (pp. 1724–1734). Association for Computational Linguistics. <https://doi.org/10.3115/v1/D14-1179>



6. Cohen, J. (1960). A coefficient of agreement for nominal scales. *Educational and Psychological Measurement*, 20(1), 37–46. <https://doi.org/10.1177/001316446002000104>
7. Congalton, R. G., & Green, K. (2019). *Assessing the accuracy of remotely sensed data: Principles and practices* (3rd ed.). CRC Press.
8. Cortes, C., & Vapnik, V. (1995). Support-vector networks. *Machine Learning*, 20, 273–297. <https://doi.org/10.1007/BF00994018>
9. Foody, G. M. (2002). Status of land cover classification accuracy assessment. *Remote Sensing of Environment*, 80(1), 185–201. [https://doi.org/10.1016/S0034-4257\(01\)00295-4](https://doi.org/10.1016/S0034-4257(01)00295-4)
10. Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep learning*. MIT Press.
11. Hinton, G. E., & Salakhutdinov, R. R. (2006). Reducing the dimensionality of data with neural networks. *Science*, 313(5786), 504–507. <https://doi.org/10.1126/science.1127647>
12. Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. *Neural Computation*, 9(8), 1735–1780. <https://doi.org/10.1162/neco.1997.9.8.1735>
13. Jolliffe, I. T., & Cadima, J. (2016). Principal component analysis: A review and recent developments. *Philosophical Transactions of the Royal Society A*, 374(2065), Article 20150202. <https://doi.org/10.1098/rsta.2015.0202>
14. Keshava, N., & Mustard, J. F. (2002). Spectral unmixing. *IEEE Signal Processing Magazine*, 19(1), 44–57. <https://doi.org/10.1109/79.974727>
15. Kumar, V., Singh, R. S., Rambabu, M., & Dua, Y. (2024). Deep learning for hyperspectral image classification: A survey. *Computer Science Review*, 53, Article 100658. <https://doi.org/10.1016/j.cosrev.2024.100658>
16. Landgrebe, D. A. (2002). Hyperspectral image data analysis. *IEEE Signal Processing Magazine*, 19(1), 17–28. <https://doi.org/10.1109/79.974718>
17. LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, 521, 436–444. <https://doi.org/10.1038/nature14539>
18. Li, S., Song, W., Fang, L., Chen, Y., Ghamisi, P., & Benediktsson, J. A. (2019). Deep learning for hyperspectral image classification: An overview. *IEEE Transactions on Geoscience and Remote Sensing*, 57(9), 6690–6709. <https://doi.org/10.1109/TGRS.2019.2907932>
19. Ma, X., Wang, H., & Geng, J. (2016). Spectral–spatial classification of hyperspectral image based on deep auto-encoder. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 9(9), 4073–4085. <https://doi.org/10.1109/JSTARS.2016.2517204>
20. Paoletti, M. E., Haut, J. M., Plaza, J., & Plaza, A. (2019). Deep learning classifiers for hyperspectral imaging: A review. *ISPRS Journal of Photogrammetry and Remote Sensing*, 158, 279–317. <https://doi.org/10.1016/j.isprsjprs.2019.09.006>
21. Pathak, D. K., Kalita, S. K., & Bhattacharya, D. K. (2021). Hyperspectral image classification using support vector machine: A spectral spatial feature based approach.



- Evolutionary Intelligence, 14, 1809–1823. <https://doi.org/10.1007/s12065-021-00591-0>
22. Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., & Duchesnay, É. (2011). Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12, 2825–2830.
 23. Roy, S. K., Krishna, G., Dubey, S. R., & Chaudhuri, B. B. (2020). HybridSN: Exploring 3-D–2-D CNN feature hierarchy for hyperspectral image classification. *IEEE Geoscience and Remote Sensing Letters*, 17(2), 277–281. <https://doi.org/10.1109/LGRS.2019.2918719>
 24. Senthil Pandi, S., Reshmy, A. K., Elangovan, D., et al. (2025). AI-driven environmental monitoring for hydroponic agriculture: ExCNN-LFCP approach. *Earth Science Informatics*, 18, Article 73. <https://doi.org/10.1007/s12145-024-01516-y>
 25. Sokolova, M., & Lapalme, G. (2009). A systematic analysis of performance measures for classification tasks. *Information Processing & Management*, 45(4), 427–437. <https://doi.org/10.1016/j.ipm.2009.03.002>
 26. Somers, B., Asner, G. P., Tits, L., & Coppin, P. (2011). Endmember variability in spectral mixture analysis: A review. *Remote Sensing of Environment*, 115(7), 1603–1616. <https://doi.org/10.1016/j.rse.2011.03.003>
 27. Vapnik, V. N. (1998). *Statistical learning theory*. Wiley.
 28. Zhu, X. X., Tuia, D., Mou, L., Xia, G.-S., Zhang, L., Xu, F., & Fraundorfer, F. (2017). Deep learning in remote sensing: A comprehensive review and list of resources. *IEEE Geoscience and Remote Sensing Magazine*, 5(4), 8–36. <https://doi.org/10.1109/MGRS.2017.2762307>
 29. Uma Rani, P., & Padmalosani. (2019). Technology adoption by farmers for the sustainability of agriculture sector. *International Journal of Recent Technology and Engineering*, 8(2S8), August 2019. <https://doi.org/10.35940/ijrte.B1094.0882S819>