PRIOR-GUIDED SPECTRAL GATING FOR COMPACT HYPERSPECTRAL RECOGNITION OF GENOME-EDITED RICE

Chao Wang[1], *Yuki Sakai*[1], *Daisuke Miyazaki*[1], *Takahiro Okabe*[2]

[1]Department of Intelligent Systems, Hiroshima City University, Japan
[2]Information Technology Track, Faculty of Engineering, Okayama University, Japan

ABSTRACT

Genome-edited rice lines can remain visually similar despite biologically meaningful differences, making nondestructive grain recognition a fine-grained spectral learning problem. Compact hyperspectral recognition requires wavelength subsets that preserve complementary information for distinguishing rice lines and milling states. We propose PBGate, a prior-guided spectral gating framework for hyperspectral recognition of genome-edited γ-aminobutyric acid (GABA)-related rice. Training-derived band-selection priors initialize wavelength-wise gates and regularize their gains during full-spectrum learning with MambaHSI. The recognizer is then further trained on the K bands specified by the prior. Experiments on ten rice line–state classes show that soft spectral gating improves grain-level overall accuracy from 89.58% to 96.88%, while PBGate achieves 89.58% using five of 51 bands and 82.29% using three bands. With the same retained bands as the corresponding standalone selectors, PBGate attains the highest grain-level accuracy among the evaluated compact methods at both band counts. The selected wavelength subsets provide candidate measurements for nondestructive multispectral screening of breeding materials.

***Index Terms*—** Hyperspectral imaging, band selection, spectral gating, rice recognition, smart agriculture

1. INTRODUCTION

Hyperspectral imaging (HSI) captures wavelength-resolved spectral–spatial signatures that enrich visual recognition [1]. In agricultural phenotyping, this information offers an opportunity to characterize biologically meaningful variation among visually similar seeds. Genome-edited γ-aminobutyric acid (GABA)-related rice provides a compelling example. GABA is a non-protein amino acid associated with plant stress responses and nutritional functionality [2]. Modifying rice glutamate decarboxylase (OsGAD) activity through calmodulin-binding-domain truncation can enhance GABA accumulation and alter stress-response phenotypes [3, 4, 5]. Recognizing these lines from intact grains could enable nondestructive screening of breeding materials while preserving samples for subsequent analysis.

This work was supported in part by JSPS KAKENHI under Grant Numbers JP26K21254, JP25K03142, and JP26K02932. The authors would like to thank Prof. Kazuhito Akama of Shimane University for providing the samples used in this study.

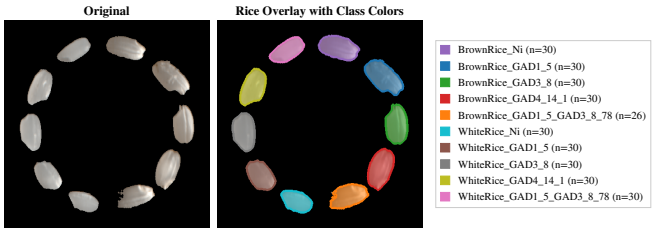


Fig. 1. Hyperspectral rice recognition across ten line–state classes. Intact grains from the OsGAD-related lines exhibit subtle appearance differences. The class-color overlay identifies the corresponding grain labels in brown and polished states.

Rice-image analysis commonly exploits differences in shape, color, and texture [6]. Here, we investigate a finer setting involving OsGAD-related rice lines and brown and polished milling states (Fig. 1). Each intact grain is assigned a line–state identity from its spectral–spatial measurements. Nondestructive recognition can preserve valuable seeds and guide sample screening. Molecular characterization [7] and biochemical profiling [2, 4] provide complementary biological information; their future integration with optical measurements could support richer genotype–phenotype analysis and breeding-oriented sample selection.

RGB integrates spectra into three broadband responses, whereas HSI preserves a densely sampled wavelength axis. Selecting a compact subset from this axis requires balancing discriminative relevance with spectral complementarity. Each retained wavelength is both a classifier input and a candidate measurement for a multispectral sensor. Band selection should therefore preserve information useful for recognition across the selected wavelengths.

Task-driven sensing couples measurement design with inference through coded illumination for reflectance-based [8] and fluorescence-based classification [9], as well as learned spectral filters [10]. Spectral–spatial CNNs [11], transformers [12], and state-space models [13] model spectral and spatial

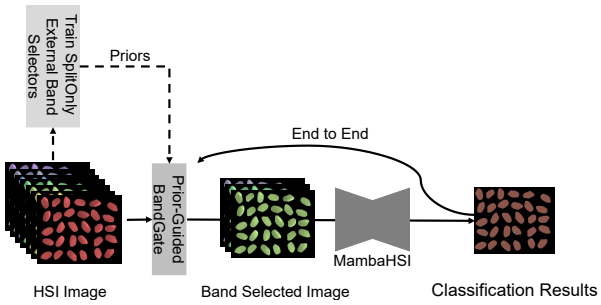


Fig. 2. PBGate framework. Training-derived wavelength priors initialize and regularize a spectral gate jointly optimized with MambaHSI. Pruning retains the prior-selected K bands for continued recognition training. The external selectors supply fixed priors.

context; DINOv3 [14] provides self-supervised dense RGB representations. We use MambaHSI [15] for spectral comparisons and DINOv3 as a pretrained RGB baseline. BandGate learns task-dependent wavelength recalibration, whose full-band and hard-pruned variants probe the distinction between continuous weighting and discrete selection: high-gain wavelengths may share information or depend on neighboring bands.

Existing selectors encode complementary spectral structure: BS-Net learns reconstruction-preserving band attention [16], NNCV uses a supervised nearest-neighbor proxy [17], and OCF identifies representative wavelengths through clustering [18]. Gumbel-Softmax [19] supports differentiable categorical learning, while Gumbel-Top-K [20] samples without replacement. These approaches provide structured wavelength candidates or task-driven discrete selection. PBGate uses selector-derived priors to guide full-spectrum gate–recognizer learning before continuing training on the prior-selected K bands.

We introduce PBGate, a prior-guided BandGate that integrates training-derived selector priors with joint gate–recognizer optimization. Prior initialization sets the wavelength gains, and prior-aware sparsity reduces the penalty on prior bands during pre-pruning training. Warm-up anchoring constrains early gain updates and is then released. Pruning retains the prior-selected K bands for continued MambaHSI training, adapting the learned representation to the compact input.

Our main contributions are threefold. First, we establish a grain-disjoint hyperspectral benchmark connecting OsGAD-related rice diversity with fine-grained optical phenotyping. Second, we introduce PBGate, a training framework that uses band-selection priors to guide full-spectrum gate–recognizer learning and subsequent adaptation to the prior-selected bands. Third, we characterize the relationship between full-band recalibration and compact-band recognition, demonstrating the highest grain-level accuracy among the evaluated compact methods at both three- and five-band budgets.

2. METHOD

PBGate combines a prior-guided spectral front end with MambaHSI [15] (Fig. 2). Let $\mathbf{X} \in \mathbb{R}^{H\times W\times B}$ be a calibrated hyperspectral image or patch, $\mathbf{x}_i \in \mathbb{R}^B$ the spectrum at foreground pixel i, and $y \in \{1,\dots,C\}$ the grain label. A compact representation retains the wavelength indices $\mathcal{S}_K \subset \{1,\dots,B\}$, with $|\mathcal{S}_K| = K$, yielding $\mathbf{X}_{\mathcal{S}_K} \in \mathbb{R}^{H\times W\times K}$.

2.1. Task-Driven Spectral Recalibration

A global BandGate learns a wavelength-dependent gain before the backbone:

$$g_b = g_{\max}\sigma(a_b/\tau), \qquad \tilde{x}_{i,b} = g_b x_{i,b}, \tag{1}$$

where a_b is a learnable logit, τ is the temperature, and $g_{\max}$ bounds the gain. Full-band training minimizes

$$\mathcal{L}_{\text{soft}} = \mathcal{L}_{\text{cls}}\big(f(\tilde{\mathbf{X}}), y\big) + \frac{\lambda_g}{B}\sum_{b=1}^{B} g_b. \tag{2}$$

The gate and classifier f are optimized jointly, allowing spectral recalibration to follow the recognition objective. Soft BandGate retains all B measurements. Its hard-pruned counterpart retains $\mathcal{S}_K^{\text{gate}} = \text{TopK}_b(g_b)$ and continues training on that subset, providing a direct comparison between full-band gains and compact-band utility.

2.2. Train-Only Spectral Priors

External selectors provide a prior set $\mathcal{P}_{K_p} \subseteq \{1,\dots,B\}$ of $K_p \le K$ distinct band indices and its mask $p_b = \mathbf{1}[b \in \mathcal{P}_{K_p}]$. We use BS-Net [16], NNCV [17], and OCF [18] to obtain reconstruction-based, supervised-proxy, and clustering-based priors, respectively. All prior-generation operations, including normalization, PCA, graph construction, nearest-neighbor scoring, reconstruction, and clustering, use training grains exclusively.

2.3. Prior-Guided BandGate

PBGate incorporates the prior into initialization, regularization, and pruning. For $0 < g_{\text{lo}} < g_{\text{hi}} < g_{\max}$, the initial gains and logits are

$$\begin{aligned} g_b^0 &= p_b g_{\text{hi}} + (1-p_b)g_{\text{lo}}, \\ a_b^0 &= \tau\,\text{logit}(g_b^0/g_{\max}). \end{aligned} \tag{3}$$

This initialization emphasizes prior-supported wavelengths while retaining learnable responses across the spectrum. The



Fig. 3. Hyperspectral acquisition setup (left) and rice examples (right). Brown rice appears above polished rice. Grains are imaged on a low-reflectance black background with randomized orientations.

pre-pruning objective combines classification with prior-aware sparsity and warm-up anchoring:

$$\mathcal{L}_{\mathrm{PBGate}} = \mathcal{L}_{\mathrm{cls}}\big(f(\tilde{\mathbf{X}}), y\big) + \frac{\lambda_g}{B}\sum_{b=1}^{B}\omega_b g_b + \frac{\lambda_a(t)}{B}\sum_{b=1}^{B}(g_b - g_b^0)^2, \tag{4}$$

where $\omega_b = rp_b + (1-p_b)$ and $0 < r < 1$. The smaller penalty favors prior bands as task-dependent gains evolve. The anchor follows

$$\lambda_a(t) = \begin{cases} \lambda_a, & t \le T_{\mathrm{warm}}, \\ 0, & t > T_{\mathrm{warm}}, \end{cases} \tag{5}$$

so that the anchor constrains early gain updates, while classification and prior-aware sparsity remain active after warm-up until pruning.

Prior-guided K-band pruning retains exactly K distinct wavelength indices. In *prior-only* selection, $K_p = K$ and $\mathcal{S}_K = \mathcal{P}_K$. In *prior-fill* selection, all prior bands are retained and the remaining $K - K_p$ indices are selected by their learned gains:

$$\mathcal{S}_K = \mathcal{P}_{K_p} \cup \operatorname*{Top}_{K-K_p,\, b\notin\mathcal{P}_{K_p}}(g_b).$$

Here Top_q returns the indices of the q largest gains over the indicated candidate set, with $\mathrm{Top}_0 = \varnothing$. Thus, when $K_p = K$, prior-fill also retains the entire prior set. The *learned* variant ranks all bands by gate scores alone. After pruning, the subset is fixed and MambaHSI continues training on $\mathbf{X}_{\mathcal{S}_K}$.

For grain g, the probabilities of its sampled patches $\mathcal{P}_g$ are aggregated as

$$\mathbf{p}_g = \frac{1}{|\mathcal{P}_g|}\sum_{p\in\mathcal{P}_g}\mathrm{softmax}\big(f(\mathbf{X}_{g,p})\big), \qquad \hat{y}_g = \arg\max_c p_{g,c}. \tag{6}$$

Each patch uses the channels retained by its sensing configuration. All compact methods share the MambaHSI architecture and this grain-level aggregation rule.

3. EXPERIMENTS

3.1. Dataset and Acquisition Protocol

The benchmark comprises 296 grains from five OsGAD-related rice lines in brown and polished states, yielding ten classes. The lines are wild-type Nipponbare (Ni), OsGAD1ΔC #5, OsGAD3ΔC #8, OsGAD4 CaMBD truncation #14–1, and a #5×#8-derived hybrid #78. Each class contains 30 grains except the brown hybrid class, which contains 26.

Hyperspectral images are acquired with an NH1H camera under artificial sunlight illumination (Fig. 3). The experiments use $B = 51$ bands spanning 450–700 nm at 5 nm intervals. Grain orientations are randomized on a low-reflectance black background, and each segmented grain forms one object for splitting and evaluation. Dark and white references yield the calibrated apparent reflectance

$$R_{\mathrm{app}}(\lambda) = \frac{I(\lambda) - I_{\mathrm{dark}}(\lambda)}{I_{\mathrm{white}}(\lambda) - I_{\mathrm{dark}}(\lambda)}. \tag{7}$$

This signal captures the combined effects of surface reflection, subsurface scattering, pose, and shading in curved grains, providing a common radiometric basis for all sensing configurations.

The training/validation/test counts are $15/5/10$ for each 30-grain class and $15/5/6$ for the 26-grain brown hybrid class, giving 150 training, 50 validation, and 96 test grains in total. The same grain-disjoint partitions are used throughout, so every patch from a given grain belongs to a single partition. RGB inputs are synthesized by integrating the calibrated cubes with Canon EOS 5D Mark II spectral sensitivities [21] under a constant-illuminant assumption. This linear broadband projection shares grains, masks, calibration, and partitions with the hyperspectral inputs, enabling a controlled comparison of spectral measurements.

3.2. Training Protocol, Baselines, and Metrics

We compare synthesized RGB, full-band MambaHSI, soft BandGate, hard-pruned BandGate, BS-Net (Conv/FC), NNCV, OCF (NC-OC-IE/MVPCA/TRC), Gumbel selection, and PBGate. We additionally evaluate pretrained DINOv3 ViT-B/16 with local linear RGB classification [14], and straight-through exact-K Gumbel selection [19, 20] with MambaHSI. Full-band and soft-gated models retain all 51 measurements; compact models use $K \in \{3, 5\}$ bands. The MambaHSI reference, two-stage, and BandGate configurations use 13×13 foreground-centered patches, with 200 sampled patches per training grain and 400 per evaluation grain. MambaHSI is trained for 100 epochs using Adam, a learning rate of 10^{-3}, and batch size 64. Hard-pruned BandGate warms up for 50 epochs, retains the Top-K gains, and continues training with the classification objective. PBGate uses the same downstream training budget and the operations in Sec. 2.3. We evaluate two pre-specified configurations: NNCV priors with prior-only

Table 1. Hyperspectral rice recognition on shared grain-disjoint partitions. DINOv3 uses local linear RGB classification; all other configurations use MambaHSI. Pixel OA excludes background pixels. Bold denotes the best value within each group.

Method / variant	Input	Grain OA	Grain AA	κ	Macro-F1	Patch OA	Pixel OA
Full-spectrum and RGB references							
Full HSI	51	0.8958	0.9000	0.8843	0.8894	0.7845	0.8589
Soft BandGate	51	**0.9688**	**0.9633**	**0.9652**	**0.9673**	**0.8136**	**0.8809**
RGB MambaHSI	3 ch.	0.7083	0.7000	0.6750	0.7016	0.5207	0.6762
RGB DINOv3 ViT-B/16 [14]	3 ch.	0.5938	0.5967	0.5483	0.5594	0.4298	0.5554
Five-band configurations ($K = 5$)							
Naive hard BandGate	5	0.7083	0.7000	0.6750	0.6956	0.5418	0.6903
BS-Net–Conv [16]	5	0.6979	0.6967	0.6636	0.6829	0.5772	0.7174
BS-Net–FC [16]	5	0.7604	0.7700	0.7336	0.7484	0.6414	0.7464
NNCV [17]	5	0.8333	0.8400	0.8146	0.8382	0.6939	0.7933
OCF–NC-OC-IE [18]	5	0.7917	0.7867	0.7680	0.7866	0.6084	0.7498
OCF–MVPCA [18]	5	0.8750	0.8733	0.8608	0.8760	**0.7252**	0.8196
OCF–TRC [18]	5	0.8229	0.8300	0.8030	0.8285	0.6683	0.7895
Gumbel Top-K [19, 20]	5	0.5000	0.4867	0.4427	0.4709	0.4160	0.5978
PBGate (ours)	5	**0.8958**	**0.9000**	**0.8841**	**0.8995**	0.7100	**0.8199**
Three-band configurations ($K = 3$)							
Naive hard BandGate	3	0.7604	0.7700	0.7336	0.7479	0.5596	0.6820
BS-Net–Conv [16]	3	0.5208	0.5333	0.4680	0.5058	0.3921	0.5819
BS-Net–FC [16]	3	0.7812	0.7900	0.7569	0.7634	0.6390	0.6990
NNCV [17]	3	0.8125	0.8200	0.7914	0.8058	0.6610	0.7595
OCF–NC-OC-IE [18]	3	0.6458	0.6400	0.6060	0.6288	0.3996	0.5725
OCF–MVPCA [18]	3	0.7708	0.7733	0.7449	0.7671	0.5864	0.7272
OCF–TRC [18]	3	0.6979	0.7033	0.6643	0.6905	0.5183	0.6839
Gumbel Top-K [19, 20]	3	0.7708	0.7800	0.7452	0.7644	0.5851	0.7174
PBGate (ours)	3	**0.8229**	**0.8300**	**0.8030**	**0.8239**	**0.6708**	**0.7730**

selection at $K_p = K = 3$, and OCF-MVPCA priors with prior-fill selection at $K_p = K = 5$. Both retain $\mathcal{S}_K = \mathcal{P}_K$; the five-band prior-fill configuration has $K - K_p = 0$ and therefore preserves the complete prior subset. Learned gains shape pre-pruning optimization in both settings.

The primary metric is grain-level overall accuracy (Grain OA), computed from Eq. (6). Grain AA, Cohen's κ, and Macro-F1 characterize class-wise recognition, while Patch OA and foreground-masked Pixel OA provide complementary spatial measures.

3.3. Spectral Recalibration and Compact Recognition

Table 1 presents the quantitative comparison. Full-band MambaHSI achieves 0.8958 Grain OA, compared with 0.7083 for RGB MambaHSI and 0.5938 for DINOv3 with local linear RGB classification. Soft BandGate further raises Grain OA to 0.9688 and Macro-F1 to 0.9673. Full-spectrum wavelength recalibration therefore improves grain-level recognition within the same MambaHSI architecture.

The hard-pruned variants attain 0.7083 and 0.7604 Grain OA at $K = 5$ and $K = 3$, respectively. The three-band result exceeds the five-band result, indicating that the retained wavelength combination matters alongside the number of bands.

Gumbel selection attains 0.7708 and 0.5000 Grain OA with three and five bands, respectively. PBGate raises Grain OA by 5.21 and 39.58 percentage points at the corresponding band counts.

At $K = 5$, PBGate obtains 0.8958 Grain OA and 0.8995 Macro-F1, compared with 0.8750 and 0.8760 for OCF-MVPCA. Its five-band Grain OA equals the value observed for ungated 51-band MambaHSI. PBGate leads the five-band comparison in Grain OA, Grain AA, κ, Macro-F1, and Pixel OA; OCF-MVPCA has the highest Patch OA of 0.7252.

At $K = 3$, PBGate reaches 0.8229 Grain OA and 0.8239 Macro-F1, compared with 0.8125 and 0.8058 for NNCV. It attains the highest values for all six metrics among the evaluated three-band methods. Compared with RGB MambaHSI, it improves Grain OA by 11.46 percentage points while using the same number of input channels.

3.4. Selected Wavelengths and Spatial Consistency

At $K = 5$, PBGate retains the OCF-MVPCA subset of 450, 495, 570, 605, and 680 nm, distributing measurements across the acquired spectral range. At $K = 3$, PBGate retains the NNCV subset of 460, 500, and 515 nm, concentrating on a smaller spectral interval. PBGate improves Grain OA while retaining the same bands as the corresponding selector. These

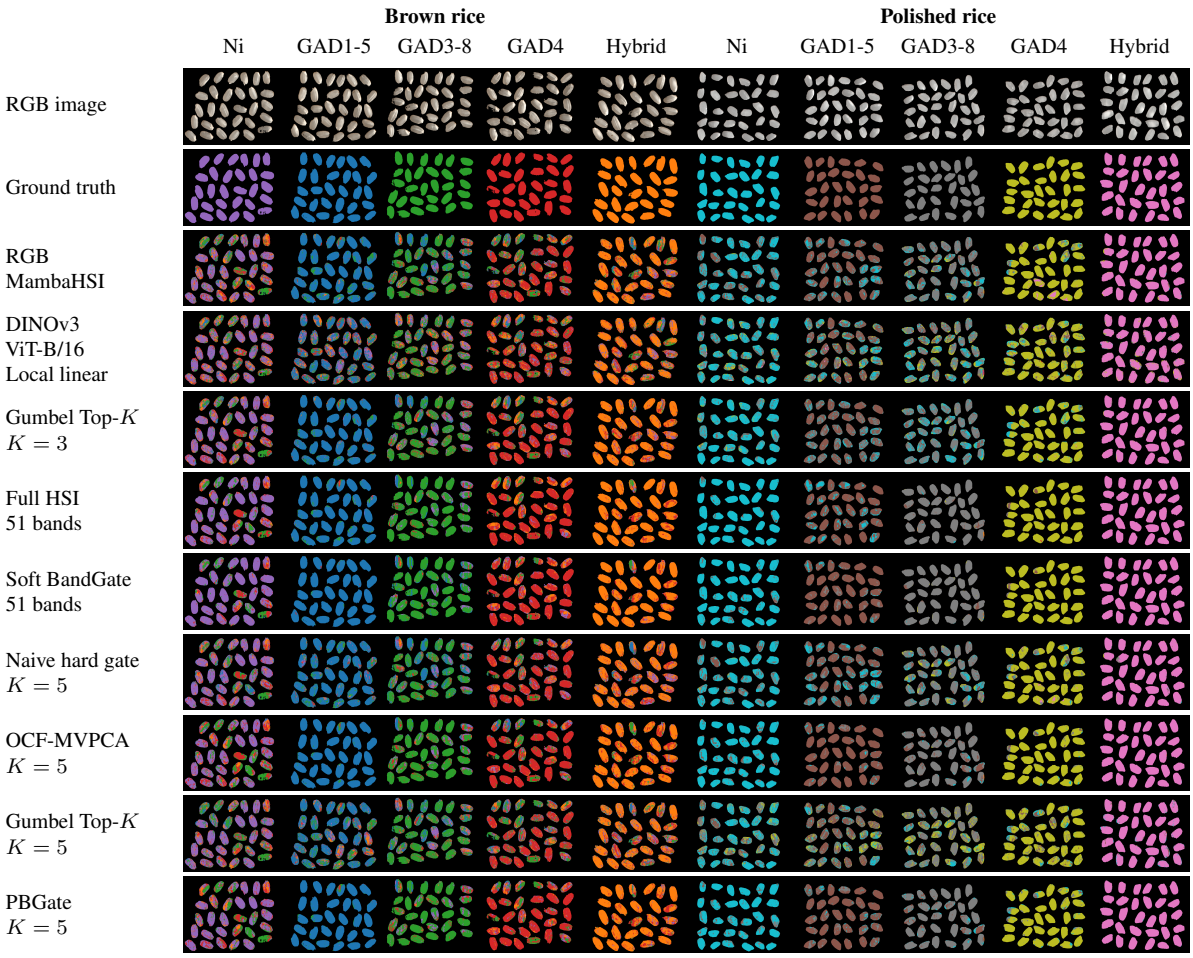


Fig. 4. Dense foreground predictions across ten line–state classes. GAD4 denotes #14–1 and Hybrid #78. DINOv3 uses local linear RGB classification; Gumbel is evaluated at $K = 3$ and 5, and the other compact methods use five bands.

matched-subset comparisons associate the improved recognition with prior-guided representation learning and subsequent adaptation to the retained bands.

Fig. 4 compares dense predictions across the ten classes, including DINOv3 and both Gumbel budgets. Full-band and soft-gated maps exhibit coherent within-grain regions in the illustrated examples. With five bands, PBGate and OCF-MVPCA attain similar foreground Pixel OA (0.8199 and 0.8196), while PBGate achieves higher grain-level accuracy. Grain-level prediction aggregates patch probabilities, so the final object decision is evaluated separately from dense per-pixel agreement.

The three- and five-band subsets provide candidate wavelength centers for multispectral acquisition. Studying their response to sensor bandwidth and illumination changes would clarify how these discrete selections transfer to physical measurements.

4. CONCLUSION

We presented PBGate for compact hyperspectral recognition of genome-edited GABA-related rice. Training-derived wavelength priors initialize and regularize spectral gates during full-spectrum learning, followed by continued MambaHSI training on the prior-selected K bands. PBGate attains the highest Grain OA among the evaluated three- and five-band methods; the five-band result matches the observed Grain OA of full-band MambaHSI. The matched-subset comparisons show that prior-guided training can improve recognition with the same retained wavelengths. Future work will evaluate multispectral acquisition with realistic sensor bandwidths and pair optical measurements with molecular and biochemical assays to study line-associated phenotypes in larger breeding collections.

5. REFERENCES

[1] M. E. Paoletti, J. M. Haut, J. Plaza, and A. Plaza, "Deep learning classifiers for hyperspectral imaging: A review," *ISPRS Journal of Photogrammetry and Remote Sensing*, vol. 158, pp. 279–317, 2019.

[2] Kazuhito Akama, Junko Kanetou, Shunsuke Shimosaki, Kohei Kawakami, Satoru Tsuchikura, and Fumio Takaiwa, "Seed-specific expression of truncated OsGAD2 produces GABA-enriched rice grains that influence a decrease in blood pressure in spontaneously hypertensive rats," *Transgenic Research*, vol. 18, no. 6, pp. 865–876, 2009.

[3] Kazuhito Akama, Nadia Akter, Hinako Endo, Masako Kanesaki, Masaki Endo, and Seiichi Toki, "An in vivo targeted deletion of the calmodulin-binding domain from rice glutamate decarboxylase 3 (OsGAD3) increases γ-aminobutyric acid content in grains," *Rice*, vol. 13, no. 1, pp. 20, 2020.

[4] Nadia Akter, Ummey Kulsum, Mohammad Moniruzzaman, Norito Yasuda, and Kazuhito Akama, "Truncation of the calmodulin binding domain in rice glutamate decarboxylase 4 (OsGAD4) leads to accumulation of γ-aminobutyric acid and confers abiotic stress tolerance in rice seedlings," *Molecular Breeding*, vol. 44, no. 3, pp. 21, 2024.

[5] Ummey Kulsum, Nadia Akter, and Kazuhito Akama, "Double-truncated version of OsGADs leads to higher GABA accumulation and stronger stress tolerance in Oryza sativa l. var. japonica," *Plant Cell Reports*, vol. 44, no. 5, pp. 95, 2025.

[6] Murat Koklu, Ilkay Cinar, and Yusuf Sinan Taspinar, "Classification of rice varieties with deep learning methods," *Computers and Electronics in Agriculture*, vol. 187, pp. 106285, 2021.

[7] Marie-Alice Fraiture, Philippe Herman, Isabel Taverniers, Marc De Loose, Dieter Deforce, and Nancy H. Roosens, "Current and new approaches in GMO detection: Challenges and solutions," *BioMed Research International*, vol. 2015, pp. 392872, 2015.

[8] Chao Wang and Takahiro Okabe, "Joint optimization of coded illumination and grayscale conversion for one-shot raw material classification," in *Proceedings of the British Machine Vision Conference (BMVC)*. 2017, pp. 136.1–136.12, BMVA Press.

[9] Yuta Asano, Misaki Meguro, Chao Wang, Antony Lam, Yinqiang Zheng, Takahiro Okabe, and Imari Sato, "Coded illumination and imaging for fluorescence based classification," in *Computer Vision – ECCV 2018*. 2018, pp. 511–526, Springer.

[10] Shijie Nie, Lin Gu, Yinqiang Zheng, Antony Lam, Nobutaka Ono, and Imari Sato, "Deeply learned filter response functions for hyperspectral reconstruction," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2018, pp. 4767–4776.

[11] Swalpa Kumar Roy, Gopal Krishna, Shiv Ram Dubey, and Bidyut B. Chaudhuri, "HybridSN: Exploring 3-d–2-d CNN feature hierarchy for hyperspectral image classification," *IEEE Geoscience and Remote Sensing Letters*, vol. 17, no. 2, pp. 277–281, 2020.

[12] Danfeng Hong, Zhu Han, Jing Yao, Lianru Gao, Bing Zhang, Antonio Plaza, and Jocelyn Chanussot, "Spectralformer: Rethinking hyperspectral image classification with transformers," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 60, pp. 1–15, 2022.

[13] Albert Gu and Tri Dao, "Mamba: Linear-time sequence modeling with selective state spaces," in *First Conference on Language Modeling*, 2024.

[14] Oriane Siméoni, Huy V. Vo, Maximilian Seitzer, Federico Baldassarre, Maxime Oquab, Cijo Jose, Vasil Khalidov, Marc Szafraniec, Seungeun Yi, Michaël Ramamonjisoa, Francisco Massa, Daniel Haziza, Luca Wehrstedt, Jianyuan Wang, Timothée Darcet, Théo Moutakanni, Leonel Sentana, Claire Roberts, Andrea Vedaldi, Jamie Tolan, John Brandt, Camille Couprie, Julien Mairal, Hervé Jégou, Patrick Labatut, and Piotr Bojanowski, "DINOv3," arXiv:2508.10104, 2025.

[15] Yapeng Li, Yong Luo, Lefei Zhang, Zengmao Wang, and Bo Du, "MambaHSI: Spatial–spectral mamba for hyperspectral image classification," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 62, pp. 1–16, 2024.

[16] Yaoming Cai, Xiaobo Liu, and Zhihua Cai, "BS-Nets: An end-to-end framework for band selection of hyperspectral image," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 58, no. 3, pp. 1969–1984, 2020.

[17] Tiancheng Zhi, Bernardo R. Pires, Martial Hebert, and Srinivasa G. Narasimhan, "Multispectral imaging for fine-grained recognition of powders on complex backgrounds," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 8699–8708.

[18] Qi Wang, Fahong Zhang, and Xuelong Li, "Optimal clustering framework for hyperspectral band selection," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 56, no. 10, pp. 5910–5922, 2018.

[19] Eric Jang, Shixiang Gu, and Ben Poole, "Categorical reparameterization with Gumbel-Softmax," in *International Conference on Learning Representations*, 2017.

[20] Wouter Kool, Herke Van Hoof, and Max Welling, "Stochastic beams and where to find them: The Gumbel-Top-k trick for sampling sequences without replacement," in *Proceedings of the 36th International Conference on Machine Learning*. 2019, vol. 97 of *Proceedings of Machine Learning Research*, pp. 3499–3508, PMLR.

[21] Jun Jiang, Dengyu Liu, Jinwei Gu, and Sabine Süsstrunk, "What is the space of spectral sensitivity functions for digital color cameras?," in *IEEE Workshop on Applications of Computer Vision*, 2013, pp. 168–179.