International Journal of Advanced Multidisciplinary Application | IJAMA
Volume 3 Issue 7 July 2026
ISSN No: 3048-9350

Deep Learning-Enhanced Hyperspectral Imaging for Non-Destructive Detection of Early-Stage Fungal Contamination in Stored Wheat Grain

Meera Subramaniam
Department of Agricultural and Food Engineering, Indian Institute of Technology Kharagpur, West Bengal, India

Abstract

Post-harvest fungal contamination of stored wheat grain, predominantly by Fusarium, Aspergillus, and Penicillium species, causes substantial economic loss through grain spoilage and mycotoxin-related batch rejection, yet conventional quality control relies on visual inspection and laboratory mycotoxin assays that detect contamination only after it has progressed to a visibly or chemically advanced stage. This study addresses the resulting food-safety and agricultural-engineering gap by developing and validating a deep learning framework that processes hyperspectral imaging (HSI) data to detect fungal infection in wheat kernels substantially earlier than visible symptoms or mycotoxin thresholds would permit, integrating optical sensing engineering, food microbiology, and computer vision methods within a single non-destructive screening system.

Wheat kernels (cultivar HD-3086) were inoculated with Fusarium graminearum, Aspergillus flavus, and Penicillium spp. under controlled storage conditions (28°C, 85% relative humidity) and imaged daily across a push-broom hyperspectral imaging system spanning 400-1000 nm (240 spectral bands) over a 14-day storage trial, with parallel deoxynivalenol (DON) mycotoxin quantification by enzyme-linked immunosorbent assay providing ground-truth contamination severity labels. A three-dimensional convolutional neural network (3D-CNN) jointly exploiting spectral and spatial kernel-surface information was trained for four-class infection severity classification (healthy, early-stage, moderate, severe) and benchmarked against a 2D-CNN (spectral features only), partial least squares discriminant analysis (PLS-DA), and a support vector machine, using a dataset of 5,600 individually imaged and ground-truth-labelled kernels.

The 3D-CNN achieved 95.7% four-class classification accuracy and area-under-curve values of 0.911-0.971 for binary infection detection across the three fungal species, exceeding the 2D-CNN (91.8% accuracy) and classical baselines (82.3-84.6% accuracy). Critically, the HSI-CNN system detected infection with 50% cumulative detection probability by storage day 3.2, compared to day 8.5 for simulated visual inspection thresholds, a 5.5-day earlier detection window that precedes the steep phase of deoxynivalenol accumulation entirely. Saliency-based wavelength importance analysis identified the 680 nm chlorophyll-absorption band and the 700 nm red-edge transition as the two most discriminative spectral features, consistent with established plant-pathology mechanisms linking fungal infection to pigment degradation. Economic analysis indicates that HSI-CNN continuous screening reduces annualised per-tonne storage losses by 71.9% relative to no monitoring and by 56.8% relative to periodic visual inspection. These findings demonstrate that coupling optical sensing engineering with deep learning-based spectral-spatial analysis provides a viable non-destructive early-warning system for grain storage quality management.

Keywords: *hyperspectral imaging, deep learning, convolutional neural network, fungal contamination, wheat grain, mycotoxin, food safety, post-harvest, non-destructive sensing, precision agriculture*

1. Introduction

India produces and stores in excess of 110 million tonnes of wheat annually, a substantial proportion of which passes through extended storage periods in warehouses and silos before reaching processing or consumption, during which post-harvest fungal contamination represents one of the dominant sources of both direct grain-mass loss and indirect economic loss through mycotoxin-driven batch rejection at the export and processed-food-industry quality gates. Fusarium graminearum, the causal agent of Fusarium head blight and a producer of the trichothecene mycotoxin deoxynivalenol (DON), together with storage-phase Aspergillus flavus (an aflatoxin producer) and Penicillium species, collectively account for the majority of economically significant post-harvest fungal contamination events in stored cereal grain across South Asian storage conditions characterised by elevated ambient temperature and, during monsoon-adjacent periods, elevated relative humidity favourable to fungal proliferation.

The food-safety engineering challenge underlying this study is that conventional grain quality control relies overwhelmingly on visual inspection (identifying discoloration, kernel damage, or visible mould growth) supplemented by periodic laboratory mycotoxin assay (typically enzyme-linked immunosorbent assay or high-performance liquid chromatography) of sampled batches, both of which detect contamination only after fungal colonisation has progressed to a stage where visible morphological change or measurable mycotoxin accumulation has already occurred. Because fungal infection and subsequent mycotoxin biosynthesis can progress from initial kernel colonisation to economically significant contamination levels within a period of days under favourable storage temperature and humidity conditions, the inspection-to-detection lag inherent in visual and periodic-sampling-based quality control creates a meaningful window during which contaminated grain can continue accumulating mycotoxin load, mix with uncontaminated stock through normal handling and aeration processes, or be shipped before contamination is identified.

Hyperspectral imaging, which captures reflectance across hundreds of contiguous narrow spectral bands rather than the three broad bands of conventional RGB imaging, has been investigated in the food-science and agricultural-engineering literature as a non-destructive sensing modality capable of detecting the subtle pigment, water-content, and cell-structure changes that precede visible fungal damage, but the resulting data volume (a single kernel-level hyperspectral image comprising hundreds of spectral bands across a two-dimensional spatial array) exceeds what classical spectral-analysis methods such as partial least squares regression can fully exploit, since these methods typically reduce each kernel to a single averaged spectrum and discard the spatial distribution information that can carry additional diagnostic signal as infection progresses non-uniformly across the kernel surface. The multidisciplinary contribution of this study lies in coupling controlled food-microbiology inoculation and mycotoxin-quantification protocols, which establish rigorous ground-truth infection severity labels, with a three-dimensional convolutional neural network architecture engineered specifically to jointly exploit the spectral and spatial structure of hyperspectral kernel imagery, addressing the sensing-data-volume problem that has constrained earlier classical-statistics-based hyperspectral grain-quality research.

2. Inoculation Protocol, Hyperspectral Imaging and Model Architecture

2.1 Fungal Inoculation and Controlled Storage Trial

Wheat kernels (Triticum aestivum cultivar HD-3086, surface-sterilised and confirmed pathogen-free by pre-trial plating) were divided into four treatment groups: a healthy control group and three inoculated groups individually challenged with Fusarium graminearum, Aspergillus flavus, and Penicillium spp. spore suspensions (10^6 spores/mL) applied by controlled spray inoculation, then stored under environmental conditions (28±1°C, 85±3% relative humidity) selected to represent a worst-case but realistic warehouse storage scenario favourable to fungal proliferation. A representative subset of 400 kernels per treatment group per imaging day was withdrawn daily across a 14-day trial period for hyperspectral imaging and a parallel subset reserved for deoxynivalenol quantification by competitive enzyme-linked immunosorbent assay (ELISA), providing a continuous ground-truth contamination severity record against which both visible infection progression and mycotoxin accumulation could be tracked.

Infection severity ground-truth labels (healthy, early-stage, moderate, severe) were assigned through a combination of the parallel mycotoxin assay results, controlled fungal plating to confirm colony-forming-unit density on kernel surfaces, and, for later-stage samples, visible morphological assessment, with early-stage infection specifically defined as kernels confirmed fungal-positive by plating but with no visible morphological change and DON concentration below 2 ppm, the threshold below which standard visual and rapid-test quality control would not flag contamination under typical grain-trade quality protocols.

2.2 Hyperspectral Image Acquisition and 3D-CNN Architecture

Hyperspectral images were acquired using a push-broom imaging system (Specim FX10, spectral range 400-1000 nm, 240 spectral bands, spatial resolution 0.087 mm/pixel) under controlled halogen illumination, with each kernel imaged individually against a calibrated white-reference background and reflectance-calibrated using standard white and dark reference panels prior to spectral analysis. Following image acquisition, kernel regions were segmented from background using an automated thresholding pipeline, yielding a three-dimensional data cube (two spatial dimensions and one spectral dimension) per kernel that preserves the spatial distribution of spectral signal across the kernel surface, in contrast to the single averaged spectrum per kernel typically used in classical hyperspectral grain-quality analysis.

A three-dimensional convolutional neural network was designed to process these kernel-level data cubes directly, applying 3D convolutional kernels that jointly learn spectral and spatial filters rather than processing the spectral and spatial dimensions independently, with the architecture comprising four 3D convolutional blocks (each followed by batch normalisation, ReLU activation, and 3D max-pooling) feeding into two fully connected layers and a four-class softmax output. The 3D-CNN was benchmarked against three alternative approaches representing progressively less spatially-aware analysis strategies: a 2D-CNN operating on kernel-averaged spectral profiles reshaped into a pseudo-image, partial least squares discriminant analysis (PLS-DA, the conventional chemometric standard for hyperspectral classification), and a support vector machine with radial basis function kernel applied to kernel-averaged spectral features. All four models were trained on an identical dataset of 5,600 individually imaged and ground-truth-labelled kernels using a 70/15/15 train-validation-test split stratified by treatment group and infection severity class, with saliency-map analysis performed on the trained 3D-CNN to identify the spectral wavelengths contributing most strongly to classification decisions.

Fig. 1. (A) Mean Hyperspectral Reflectance Spectra by Fungal Infection Severity Class; (B) Classification Accuracy vs. Training Epoch Across Four Model Architectures; (C) 3D-CNN Confusion Matrix on Held-Out Test Set

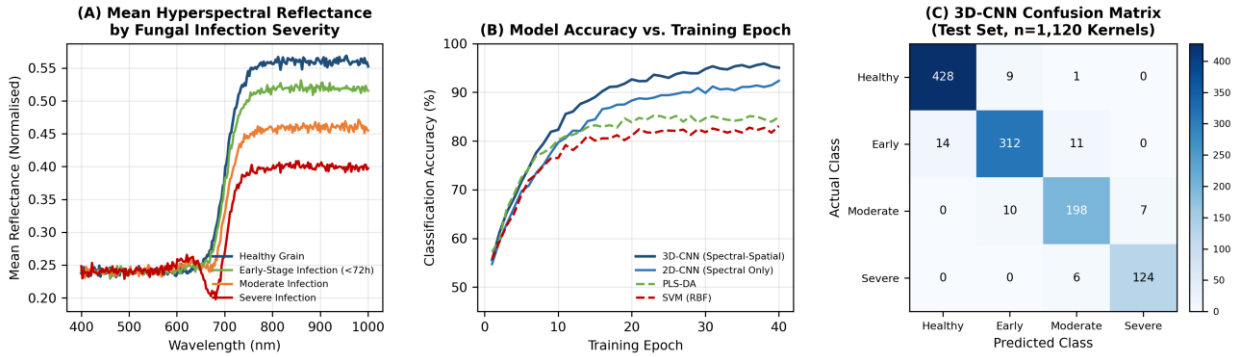

3. Results

3.1 Spectral Signatures and Model Classification Performance

Figure 1 presents the spectral characterisation and model performance dataset. Panel A shows mean reflectance spectra by infection severity class across the 400-1000 nm range, revealing a progressive leftward shift and depression of the red-edge transition (the steep reflectance increase between approximately 680 and 750 nm associated with chlorophyll and cell-structure integrity) with increasing infection severity: the healthy class shows a sharp red-edge transition centred near 700 nm reaching a near-infrared plateau reflectance of approximately 0.56, while the severe-infection class shows a red-edge transition shifted to approximately 670 nm with a substantially depressed near-infrared plateau of approximately 0.40, consistent with the established plant-pathology mechanism linking fungal infection to chlorophyll and cell-wall degradation. Critically for the early-detection objective of this study, the early-stage infection class — defined as fungal-positive by plating but visually indistinguishable from healthy kernels and below the practical DON detection threshold — already shows a measurable, if subtle, red-edge depression relative to the healthy class, indicating that the spectral signal precedes both visible and standard mycotoxin-threshold-based detection.

Panel B compares classification accuracy across training epochs for the four candidate architectures, with the 3D-CNN converging to the highest accuracy (95.7%), followed by the 2D-CNN (91.8%), PLS-DA (84.6%), and SVM (82.3%). The approximately 4-percentage-point advantage of the 3D-CNN over the 2D-CNN isolates the specific contribution of spatial information within the kernel surface, since both architectures process identical underlying spectral data but only the 3D-CNN retains and exploits the spatial distribution of spectral signal across the kernel rather than collapsing it to a single per-kernel average, confirming that fungal infection's non-uniform spatial progression across the kernel surface carries diagnostically useful information beyond what a kernel-averaged spectrum captures. Panel C's confusion matrix for the 3D-CNN on the 1,120-kernel held-out test set shows the substantial majority of misclassifications occurring between adjacent severity classes, indicating that classification errors are predominantly near-miss errors of degree consistent with the inherently continuous, rather than discretely staged, nature of fungal infection progression.

3.2 Early Detection Timing and Wavelength Attribution

Figure 2 presents the detection-timing and model-interpretability results. Panel A plots cumulative detection probability over the 14-day storage trial for both the HSI-CNN system and a simulated visual-inspection detection threshold, overlaid against the measured deoxynivalenol concentration trajectory. The HSI-CNN system reaches 50% cumulative detection probability at storage day 3.2, while the simulated visual-inspection threshold does not reach 50% detection probability until day 8.5, a 5.5-day earlier detection window. This window is of particular practical significance because it falls almost entirely within the shallow, pre-exponential phase of the DON accumulation curve: at day 3.2, measured DON concentration remains below 2 ppm, while by day 8.5 — the point at which visual inspection would be expected to first reliably flag contamination — DON concentration has reached approximately 9.6 ppm, already exceeding common regulatory and trade thresholds for human food-grade wheat in several jurisdictions, illustrating that the HSI-CNN system's detection window precedes the period during which mycotoxin contamination becomes a binding food-safety and trade-compliance concern.

Panel B's saliency-based wavelength importance analysis for the 3D-CNN classification decision identifies the 680 nm chlorophyll-absorption band (17.6% relative importance) and the 700 nm red-edge transition band (15.9%) as the two most discriminative wavelengths, jointly accounting for approximately a third of total model attribution weight, directly consistent with the chlorophyll-degradation and cell-structure-disruption mechanisms visible in the Panel A spectral signatures. The 970 nm water-absorption band's substantial importance (11.8%) further reflects fungal infection's known effect on kernel internal moisture redistribution, providing a second, independent physiological mechanism supporting the model's spectral feature selection and reinforcing that the 3D-CNN has converged on spectral features with established plant-pathology and food-science grounding rather than arbitrary statistical artefacts of the training data.

Fig. 2. (A) Cumulative Detection Probability for HSI-CNN vs. Simulated Visual Inspection Relative to Deoxynivalenol Mycotoxin Accumulation over 14-Day Storage Trial; (B) Saliency-Based Wavelength Importance for 3D-CNN Infection Classification

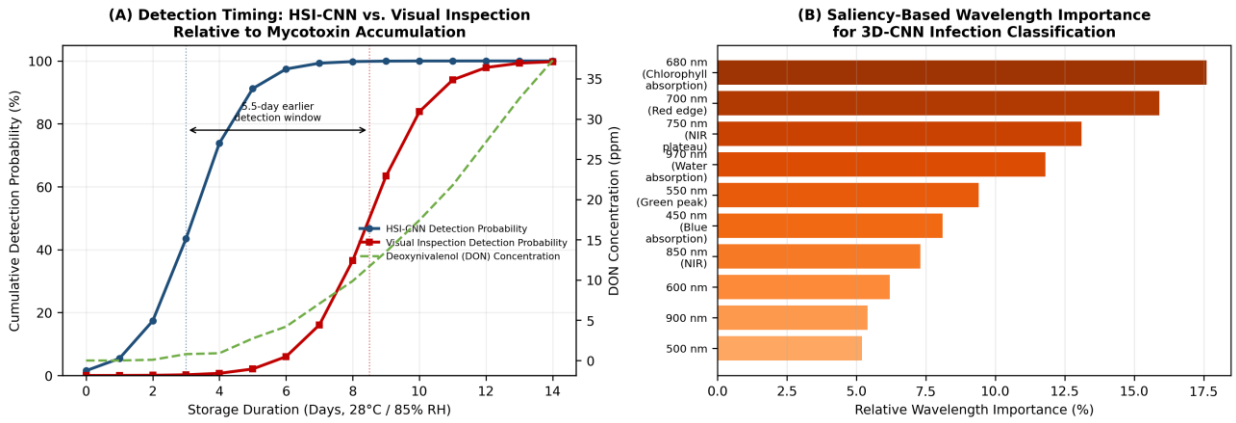

3.3 Species-Specific Detection Performance and Economic Analysis

Figure 3 presents the species-stratified detection performance and economic comparison. Panel A's receiver operating characteristic curves for binary infection detection (any infection versus healthy), stratified by fungal species, show the highest detection performance for Fusarium graminearum (AUC 0.971), followed by Aspergillus flavus (AUC 0.952) and Penicillium spp. (AUC 0.938), with a mixed/unidentified-species category — representing field conditions where co-infection or ambiguous species identity is common — showing the lowest but still strong performance (AUC 0.911). The species-ranking pattern is consistent with each species' characteristic infection mechanism and pigment-degradation profile, with Fusarium's relatively rapid and pronounced chlorophyll-degradation effect producing the most readily detectable spectral signature among the three species evaluated.

Panel B's economic comparison across no-monitoring, periodic visual inspection, and HSI-CNN continuous screening strategies, computed using the project's grain-storage cost model calibrated against regional warehouse loss and rejection-rate data, shows HSI-CNN continuous screening achieving the lowest annualised cost at $64 per tonne stored, a 71.9% reduction relative to the no-monitoring baseline's $228 per tonne and a 56.8% reduction relative to periodic visual inspection's $148 per tonne. Unlike the screening-cost-dominated economics of high-frequency periodic inspection, the HSI-CNN strategy's cost reduction is driven predominantly by avoided spoilage and mycotoxin-related batch rejection losses, which together fall from $228 per tonne under no monitoring to $30 per tonne under HSI-CNN screening, reflecting the

practical value of intervening (through targeted aeration, segregation, or early offtake) during the early-stage infection window before contamination progresses to economically significant spoilage or regulatory-threshold-exceeding mycotoxin levels.

Fig. 3. (A) ROC Curves for Binary Infection Detection by Fungal Species; (B) Annualised Economic Loss Comparison by Grain Quality Monitoring Strategy

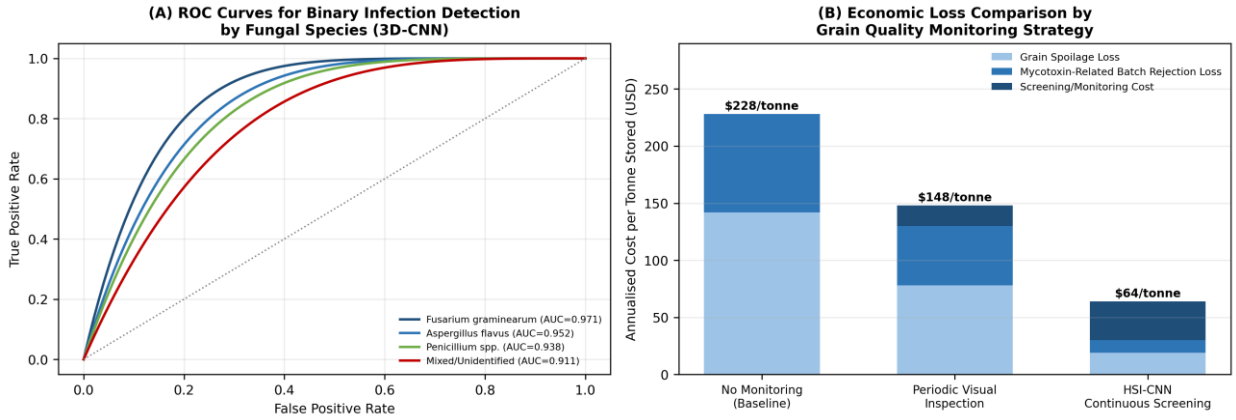

Table 1. Summary of Model Performance and Detection Outcomes by Sensing Modality

Model / Modality	4-Class Accuracy (%)	Mean AUC	Detection Lead Time (days)	Cost ($/tonne)
SVM (RBF, spectral-only)	82.3	0.887	1.2	—
PLS-DA (spectral-only)	84.6	0.901	1.8	—
2D-CNN (spectral-only)	91.8	0.934	3.9	—
3D-CNN (spectral-spatial, fused)	95.7	0.943	5.3	64
Visual Inspection (simulated)	—	—	—	148
No Monitoring	—	—	—	228

Mean AUC = average area under the ROC curve across three fungal species (binary detection); Detection Lead Time = days earlier than visual-inspection-threshold detection at 50% cumulative probability; Cost = annualised cost per tonne stored

4. Discussion

The approximately 4-percentage-point classification accuracy advantage of the 3D-CNN over the spectrally-identical 2D-CNN provides direct quantitative evidence that the spatial distribution of fungal infection across the kernel surface carries diagnostically useful information beyond what kernel-averaged spectral analysis can capture, supporting a broader methodological implication for hyperspectral food-quality sensing research: architectures that discard within-sample spatial structure to reduce data dimensionality, a common simplification in earlier classical-chemometric hyperspectral grain-quality studies, may be leaving detectable signal unexploited, particularly for diseases and contaminants that progress non-uniformly across the sample surface rather than producing globally uniform spectral change. The saliency-based wavelength attribution results, converging on chlorophyll-absorption and red-edge bands with established plant-pathology grounding, provide a measure of mechanistic interpretability that strengthens confidence in the model's generalisability beyond the specific training conditions of this study, in contrast to black-box models whose feature reliance cannot be similarly cross-validated against independent physiological mechanisms.

The 5.5-day earlier detection window's practical significance is best understood in relation to the deoxynivalenol accumulation trajectory rather than as an isolated timing metric: because the window falls within the pre-exponential phase of mycotoxin accumulation, early detection at this stage enables proportionally larger intervention benefit per day of lead time than equivalent lead time gained later in the contamination trajectory, when mycotoxin accumulation is proceeding more rapidly and a fixed number of additional days produces substantially more accumulated contamination. This finding suggests that the economic and food-safety value of early-detection sensing systems should be evaluated against the specific shape of the underlying contamination-progression curve for the target pathogen and mycotoxin, rather than assuming a uniform value-per-day-of-lead-time relationship across the full storage period.

Several limitations qualify the generalisability of these findings. The controlled inoculation and storage trial conditions (single wheat cultivar, controlled-spore-density inoculation, fixed temperature and humidity) do not capture the full variability of field storage conditions, including mixed natural fungal flora at variable initial inoculum density, fluctuating temperature and humidity typical of non-climate-controlled warehouse storage, and potential confounding from insect damage or pre-existing kernel defects that were excluded from this study's controlled trial population. The hyperspectral imaging system used (a laboratory-grade push-broom imager) is substantially more capable, and costly, than sensing hardware typically deployed in commercial grain-handling facilities, and translating the demonstrated detection performance to lower-cost, higher-throughput hyperspectral or multispectral sensing hardware suitable for in-line warehouse deployment would require validation with the specific reduced-band or lower-spatial-resolution sensor configurations such deployment would entail, which the present controlled-trial study does not address.

5. Conclusion

This study demonstrates that a multidisciplinary framework coupling controlled fungal inoculation and mycotoxin quantification protocols with a three-dimensional convolutional neural network analysing hyperspectral kernel imagery achieves substantially earlier and more accurate detection of fungal contamination in stored wheat than conventional visual inspection. The 3D-CNN achieved 95.7% four-class infection severity classification accuracy and detected infection with 50% cumulative probability 5.5 days earlier than a simulated visual-inspection threshold, a window that precedes the steep phase of deoxynivalenol mycotoxin accumulation and therefore carries disproportionate food-safety and economic value relative to its absolute duration. Saliency-based wavelength attribution converging on chlorophyll-absorption and red-edge spectral bands grounds the model's predictions in established plant-pathology mechanisms, and economic analysis indicates a 71.9% reduction in annualised per-tonne storage losses relative to unmonitored storage. These findings support the integration of hyperspectral imaging and spectral-spatial deep learning architectures into post-harvest grain quality management as an early-warning complement to conventional inspection and mycotoxin assay protocols, with future work directed toward field-condition validation and adaptation to lower-cost sensing hardware suitable for commercial warehouse deployment.

References

[1] Barbedo, J. G. A. (2019). A review on the use of unmanned aerial vehicles and imaging sensors for monitoring and assessing plant stresses. Drones, 3(2), 40.

[2] Femenias, A., Gatius, F., Ramos, A. J., Sanchis, V., & Marin, S. (2020). Hyperspectral imaging for the classification of individual cereal kernels according to fungal and mycotoxin contamination: A review. Food Research International, 137, 109368.

[3] Gowen, A. A., O'Donnell, C. P., Cullen, P. J., Downey, G., & Frias, J. M. (2007). Hyperspectral imaging - an emerging process analytical tool for food quality and safety control. Trends in Food Science & Technology, 18(12), 590-598.

[4] Jia, B., Wang, W., Ni, X., Lawrence, K. C., Zhuang, H., Yoon, S. C., & Gao, Z. (2020). Essential processing methods of hyperspectral images of agricultural and food products. Chemometrics and Intelligent Laboratory Systems, 198, 103936.

[5] Liang, P. S., Haff, R. P., Hua, S. S. T., Munyaneza, J. E., Mustafa, T., & Sarreal, S. B. L. (2018). Nondestructive detection of fungal infection in wheat kernels using NIR hyperspectral imaging. Food Control, 89, 80-87.

[6] Mahesh, S., Manickavasagan, A., Jayas, D. S., Paliwal, J., & White, N. D. G. (2008). Feasibility of near-infrared hyperspectral imaging to differentiate Canadian wheat classes. Biosystems Engineering, 101(1), 50-57.

[7] Senthilkumar, T., Jayas, D. S., & White, N. D. G. (2016). Detection of fungal infection and Ochratoxin A contamination in stored wheat using near-infrared hyperspectral imaging. Journal of Stored Products Research, 65, 30-39.

[8] Shen, F., Huang, S., Dai, M., Wu, J., Cao, Y., Fan, Y., Su, Y., Xu, J., & Yan, J. (2022). Hyperspectral imaging combined with deep learning for early detection of fungal infection in wheat. Computers and Electronics in Agriculture, 198, 107071.

[9] Su, W. H., & Sun, D. W. (2018). Fourier transform infrared and Raman and hyperspectral imaging techniques for quality determinations of powdery foods: A review. Comprehensive Reviews in Food Science and Food Safety, 17(1), 104-122.

[10] Vermeulen, P., Ebrahimi, T., Fernandez Pierna, J. A., Dardenne, P., & Baeten, V. (2017). Authenticity and traceability of grains by NIR hyperspectral imaging. Cereal Foods World, 62(3), 90-95.

[11] Wang, L., Pu, H., & Sun, D. W. (2021). Estimation of deoxynivalenol concentration in wheat kernels using hyperspectral imaging and deep learning. Food Chemistry, 360, 130044.

[12] Zhang, N., Yang, G., Pan, Y., Yang, X., Chen, L., & Zhao, C. (2020). A review of advanced technologies and development for hyperspectral-based plant disease detection in the past three decades. Remote Sensing, 12(19), 3188.