WheatScout: A Handheld Analyzer for Pre-Harvest Wheat Protein and Moisture Monitoring

Shanwen Chen[1], Haiyan Hu[1], Meng Yang[2], Qian Zhang[1*]
[1]Department of Computer Science and Engineering, The Hong Kong University of Science and Technology, HongKong, China
[2]The College of Agriculture, Shandong Agricultural University, Shandong, China
schenfh@connect.ust.hk,hhuap@connect.ust.hk,yangmeng@sdau.edu.cn,qianzh@cse.ust.hk

Abstract

Pre-harvest protein and moisture content in wheat jointly determine market value and optimal harvest timing. However, existing detection methods are limited to expensive laboratory spectrometers or handheld devices requiring threshed grains, rendering practical in-field quality assessment inaccessible to smallholder farmers. Achieving low-cost, in-situ measurement requires overcoming three critical challenges: strong water absorption in fresh spikes masking weaker protein signals, morphological variations across cultivars causing spectral baseline shifts, and scattering-related distortions introduced by variable probe contact during handheld operation. To address these issues, we present WheatScout, a co-designed system integrating optimized optics with task-aware learning. The hardware features a diffuse optical path to minimize geometric bias, while the processing pipeline employs a moisture-first framework to explicitly decouple water interference prior to protein estimation. Furthermore, we integrate domain-invariant representation learning to enhance cross-cultivar generalization and a trainable scattering correction module to compensate for residual field-trial errors. During a 15-day harvest window, we systematically collected 600 wheat ear samples, covering 10 different varieties and five functional categories, thus constructing a dataset with broad spatiotemporal representativeness and varietal diversity. On this dataset, WheatScout achieved an $R^2 > 97\%$ for moisture and an $R^2 > 92\%$ for protein, which are comparable to lab-grade spectrometers at two orders of magnitude lower cost.

CCS Concepts

• **Human-centered computing** → **Mobile devices**; • **Computing methodologies** → **Spectral methods**.

Keywords

Handheld sensing, Wheat quality monitoring, Near-infrared spectroscopy

ACM Reference Format:
Shanwen Chen[1], Haiyan Hu[1], Meng Yang[2], Qian Zhang[1]. 2026. WheatScout: A Handheld Analyzer for Pre-Harvest Wheat Protein and Moisture Monitoring. In *The 24th Annual International Conference on Mobile Systems, Applications and Services (MobiSys '26), June 21–25, 2026, Cambridge, United Kingdom.* ACM, New York, NY, USA, 13 pages. https://doi.org/10.1145/3745756.3809188

*The corresponding author is Qian Zhang (qianzh@cse.ust.hk).

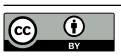

MobiSys '26, Cambridge, United Kingdom

ACM ISBN 979-8-4007-2027-7/2026/06
https://doi.org/10.1145/3745756.3809188

1 Introduction

Wheat is one of the world's most important staple foods, with a global planting area of 219.6 million hectares and a total output of approximately 731 million tons in 2021, feeding approximately 2.5 billion people [17]. Smallholder farmers, defined generally as those working less than 2 hectares, are the backbone of the global food system, producing an estimated 70-80% of its supply [7, 26, 33]. For them, harvest timing is a direct and critical variable affecting farm profitability [3], as it governs both *yield preservation* and *grain quality*. Timely harvesting can reduce yield losses by 5% to 10% and, more critically, locks in optimal grain quality by determining final moisture and protein content, which directly impact market prices [19].

These changes are concentrated within a narrow window of approximately 7-10 days from the late milk stage to full maturity, during which grain moisture content decreases by 1-3 percentage points daily, while protein content begins to decline after reaching a brief peak [8, 25, 27]. Critically, shifts as small as 0.5% in protein or 1-2% in moisture around market thresholds are sufficient to alter the payment tier for an entire harvest [3, 27]. However, these critical changes are imperceptible to farmers relying on qualitative methods like visual inspection or hand-feel, creating a fundamental data gap that precludes optimized decision-making. Existing technological solutions, however, are inadequate for in-field use by smallholders. Benchtop VNIR/NIR spectrometers, while analytically capable, are cumbersome, stationary, and require destructive sampling and laboratory analysis, creating delays that exceed the harvest window [32]. Commercial handheld devices can conduct in-field analysis but often still require threshed kernels and are too costly for broad adoption [5]. Lower-cost options lack essential capabilities: dielectric moisture meters do not measure protein [18], and vision-based or aerial sensing primarily reflects canopy-level properties rather than grain composition [29]. Therefore, an **affordable**, **field-deployable** device for **non-destructive**, quantitative analysis of protein and moisture directly on intact wheat ears is needed.

To bridge this gap, this paper proposes **WheatScout**, a hardware-algorithm co-designed system aimed at bringing laboratory-level protein and moisture analysis capabilities directly to the field, applicable to whole wheat ears, while reducing costs to a level acceptable to smallholder farmers. At the heart of **WheatScout** is an ultra-low-cost multi-LED VNIR (visible-near-infrared) sensor, costing

one to two orders of magnitude less than benchtop instruments. This sensor is encapsulated in a custom 3D-printed housing and can complete a measurement within 500 milliseconds via a single-button operation. However, this extreme cost-effectiveness presents significant technical challenges: sparse and non-uniform spectral bands, moderate optical resolution, and sensitivity to measurement geometry and lighting artifacts. Therefore, directly applying standard chemometric models developed for laboratory instruments is doomed to fail, new algorithms must be developed to compensate for the inherent limitations of the hardware.

Non-destructive field measurements of whole wheat ears must address three interconnected core challenges. **(1) Dominant water interference**: In fresh, unharvested ears, the spectral signal is dominated by strong water absorption. The variations in water content cause large, non-linear shifts in the signal that easily obscure the much weaker spectral signatures from protein. **(2) Confounding effects of plant structure**: Wheat varieties have different physical structures (e.g., husk, awn). These morphological differences change how light scatters off the ear, creating spectral patterns that are unrelated to chemical composition. A model trained on one variety may learn these "shortcut" features and perform poorly on others. **(3) Uncontrollable measurement environment**: Handheld field measurements inevitably introduce variations in probe angle, distance, and contact pressure, leading to changes in photon path length and scattering effects. This introduces significant and unpredictable noise into the spectral data, which is particularly damaging when working with sparse data points.

To address the aforementioned challenges, WheatScout does not employ a simple end-to-end regression approach, but instead adopts a phased, task-aware learning strategy. Our algorithmic framework is as follows: **First, scattering correction (Challenge (3)).** We employ a hardware-algorithm co-design, integrating physical light homogenization with multi-scale CNNs to filter baseline drift, yielding high-fidelity signals immune to geometric variations. **Second, water compensation (Challenge (1)).** We propose a "moisture-first" framework that explicitly models and subtracts nonlinear water absorption, reconstructing a normalized "dry-basis" spectrum to decouple moisture interference. **Finally, chemistry-focused regression (Challenge (2)).** We implement a domain-invariant strategy combining dual-level adversarial training with self-attention. This mechanism purges variety-specific morphological noise while capturing global protein dependencies, ensuring reliance on genuine biochemical signals.

We realize WheatScout as a compact prototype that embeds a commercial multi-LED VNIR spectrometer inside a custom 3D-printed enclosure. The device weighs under 100 g, operates via a single button, and completes each measurement in approximately 500 ms, capturing both high- and low-drive readings for all 22 LED channels. The enclosed design ensures that the ambient light level within the sample chamber is less than 1 lux.

To verify the robustness of WheatScout in real-world farmland environments, we collaborated with the College of Agriculture to conduct a systematic field trial in a typical wheat-growing region. The trial collected 600 complete wheat ear samples, covering several regionally representative winter wheat varieties and spanning the 15-day pre-harvest window from milk stage to full maturity. All samples were obtained in situ in the field, ensuring that the data are

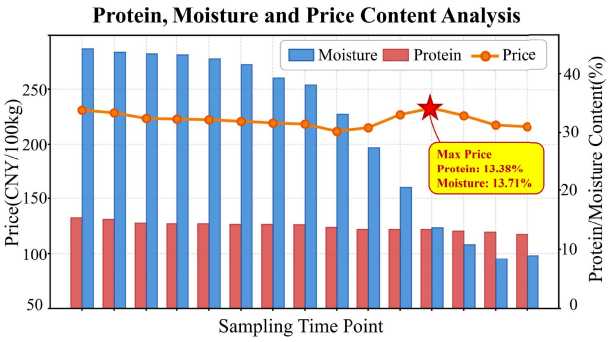

Figure 1: Moisture, protein, and resulting price trajectory for pre-harvest wheat over the 15 days preceding full maturity.

representative of real field conditions. On this extensive and representative dataset, WheatScout demonstrated superior performance, achieving moisture prediction $R^2 > 97\%$ (up 11.6%) and protein prediction $R^2 > 92\%$ (up 13.71%) over baseline, compellingly validating the system's capability to overcome structural heterogeneity and environmental uncertainty in field settings. This work makes three main contributions:

- We propose an **optical-algorithm collaborative correction method** tailored for complex field scattering environments. By integrating a fixed-geometry probe with double-sided diffusive optics and a multi-scale CNN, we effectively mitigate geometric artifacts and stabilize spectral baselines directly at the source.
- We develop a **"moisture-first" spectral decoupling framework** driven by conditional GANs and dual-domain adversarial training. This approach reconstructs latent dry-basis features from wet-ear measurements, explicitly disentangling water interference from protein signals while ensuring generalization across diverse cultivars.
- We collect a comprehensive **field wheat spectral dataset** comprising 600 samples across 10 distinct varieties and report the full system implementation of WheatScout. Extensive validation demonstrates the system's feasibility, matching laboratory-grade accuracy under extreme cost constraints.

2 Background

Pre-Harvest Wheat Quality Dynamics. The core of wheat quality monitoring lies in the accurate determination of protein and moisture, as these two parameters directly determine the market and usability of the grain. Figure 1 illustrates moisture and protein dynamics for Jimai 22, a widely grown, high-yield bread wheat cultivar in northern China. Using the harvest-and-drying economics model of Abawi [2], the resulting grain price reaches its maximum at the point marked in Figure 1, where moisture is still moderate and protein remains relatively high. This peak reflects an economically favorable harvest window in which the trade-off between drying costs and protein-related premiums is most advantageous. Missing this window will result in significant losses: harvesting too early increases drying costs due to high moisture content, while harvesting too late leads to yield losses due to grain loss and reduced quality.

Field Wheat Detection with Spectroscopy. Mainstream on-site wheat testing relies on near-infrared (NIR) spectroscopy, which measures absorption by molecular bonds such as C–H, N–H, and O–H. Different components exhibit characteristic absorption bands: for example, moisture is detected via O–H absorption, while protein involves N–H and C–H absorptions. In field measurements of whole wheat ears, the spectral signal is influenced by multiple factors, described generally as:

$$I(\lambda) = I_0(\lambda) \cdot e^{-\mu_m(\lambda) \cdot c_m \cdot l} \cdot R(\lambda) + \epsilon, \tag{1}$$

where λ is the light wavelength, $\mu_m(\lambda)$ is the absorption coefficient of component m, c_m is its concentration, l is the effective optical path, $R(\lambda)$ is the scattering effect, and ϵ is the measurement noise. It is worth noting that under field measurement conditions, factors such as fluctuations in light source, changes in measurement geometry, and fluctuations in ambient temperature introduce additional variations, significantly increasing the difficulty of signal interpretation.

Rationale of Grain Analysis with Sparse Spectra. The accurate discrimination of overlapping absorption features in practice requires high-resolution reflectance spectra $R(\lambda)$. However, practical field systems typically provide only a small number of discrete bands. To bridge this gap, spectral reconstruction can be used to map such limited observations I_{lim} to lab-like responses $I_{\text{lab}}(\lambda) = E(\lambda)R(\lambda)S(\lambda)$ by exploiting local spectral smoothness and the low-dimensional structure of reflectance spectra [15, 16, 28]. Proper constraints (non-negativity, smoothness) and explicit modeling of illumination $E(\lambda)$ and sensor sensitivity $S(\lambda)$ help stabilize baselines. For fresh wheat ears, reconstruction can recover weak absorption cues that are otherwise masked by strong water absorption near ~970/~1450 nm.

3 Data Collection and Analysis

3.1 Data Collection

Field experiments were conducted in Shandong Province, China. Shandong has led agricultural exports for 26 consecutive years, reaching a total value of 23 billion US dollars in 2024 [6], and contributes approximately 27 million tons of wheat annually, accounting for 18% of the national output. Spectral data were acquired over a 15-day pre-harvest period from 25 May to 8 June 2025, with field conditions illustrated in Figure 2. The study examined ten wheat varieties representing the genetic diversity of Northern China, spanning five distinct functional categories as detailed in Table 1. The selection features **major varieties** representing current production standards, notably **Jimai 22**, having maintained the largest annual planting area in China for the past 30 years with a cumulative cultivation exceeding 24 million hectares. To capture diverse physiological traits, the cohort includes **high-yield varieties** that push yield limits through increased spike numbers and grain weight; **high-value varieties** characterized by strong gluten and superior processing quality for high-end food production; **dryland varieties** ensuring stability in non-irrigated, water-scarce environments; and the **special variety** Lanmai, a blue-grained wheat rich in anthocyanins and antioxidants targeting functional nutrition. A comprehensive dataset was constructed to capture the spectral variability of wheat. A total of 600 samples were collected over 15

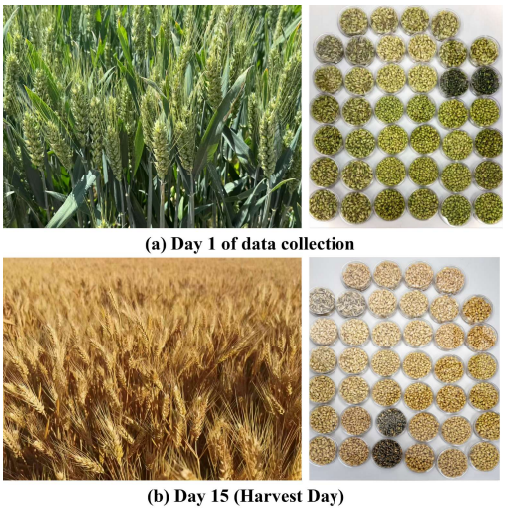

(a) Day 1 of data collection

(b) Day 15 (Harvest Day)

Figure 2: Field conditions during the 15-day pre-harvest data collection period.

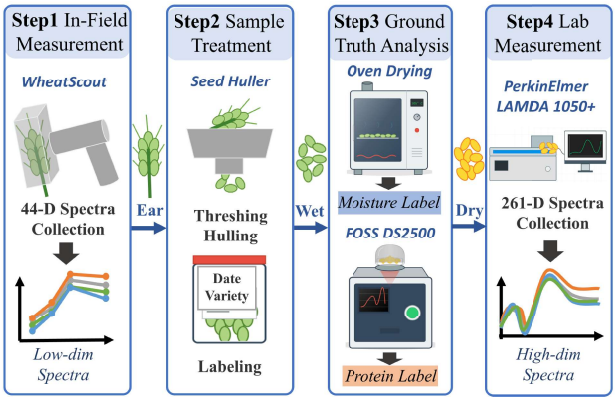

Figure 3: The overall data acquisition workflow, illustrating the transition from in-field spectral collection to laboratory analysis.

consecutive days, consisting of four randomized spikes per variety daily. As illustrated in Figure 3, the data acquisition and processing pipeline consists of four distinct stages:

- **In-Field Measurement.** Each spike was measured in situ using WheatScout, which captured 44D spectral data (200–1700 nm). To ensure data reliability, scans were averaged from the apical, middle, and basal segments to obtain a representative low-dimensional spectrum.
- **Sample Treatment.** Immediately after field collection, each wheat sample underwent threshing and hulling to isolate pure grains. The resulting wet samples were then systematically labeled with their specific collection date and variety information.
- **Ground Truth Analysis.** To establish reference values, moisture content was first determined via low-temperature (55 °C) oven drying, which prevents protein denaturation. Subsequently, the

Table 1: Wheat varieties and their characteristics

Characteristic	Variety
Major variety	Jimai 22, Shannong 38, Shannong 28
High-yield variety	Shannong 77, Shannong 73
High-value variety	Shannong 66, Shannong 47
Dryland variety	Shannong 57, SNM422
Special blue-grain variety	Lanmai

protein content of the dried grains was quantified using a FOSS DS2500 grain analyzer.

- **Laboratory Measurement.** Finally, 261D spectra (400–1700 nm) of the dried grain samples were acquired using a PerkinElmer LAMBDA 1050+ UV/Vis/NIR spectrometer, serving as the high-fidelity benchmark.

3.2 Data Analysis

To leverage the dataset's scale for robust system design, we analyzed its spectral characteristics to pinpoint physical variations our framework must resolve.

3.2.1 Dataset Characteristics. The samples exhibited a wide dynamic range of moisture (7.36%–46.85%, mean = 30.75%, SD = 13.81%) and dry-basis protein content (9.27%–15.44%, mean = 13.00%, SD = 1.11%). This protein range covers China's primary commercial wheat categories, including weak- (< 12.5%), medium- (≥12.5%), medium-strong- (≥13%), and strong-gluten (≥14%) varieties [21]. As shown in Figure 1, the final 15-day maturation phase is characterized by a rapid moisture decline contrasted with a slow decrease in protein. Beyond these statistics, the dataset captures three critical dimensions of spectral variability: moisture-induced non-linearity, cultivar-specific baselines, and geometric scattering.

3.2.2 Non-linear spectral variation under moisture fluctuations. In intact grains, moisture not only reshapes surface optical properties but also introduces water absorption bands. Higher water content plasticizes the pericarp and smooths the surface, enhancing multiple scattering and broadly elevating diffuse reflectance across the VNIR range [13]. Crucially, our analysis confirms that moisture effects extend beyond a simple baseline shift [23]. As shown in Figure 4b, within the protein-sensitive regions, particularly near the N–H(II) and C–H(II) bands, wet samples differ from dry ones not only by a distorted spectral shape but also by a significantly wider range of intensity fluctuations at any given wavelength. This occurs because water absorbs strongly around 950–1000 nm and 1400–1500 nm, so when moisture deviates from the calibration range, its spectral distortions directly interfere with protein-related features. Consequently, moisture variation is spuriously transferred into protein predictions, leading to systematic bias and large estimation errors. This mechanism is consistent with the performance comparison in Figure 4a, where a model that performs well on dry-basis spectra (R^2 = 0.95) degrades substantially on wet-basis spectra (R^2 = 0.68).

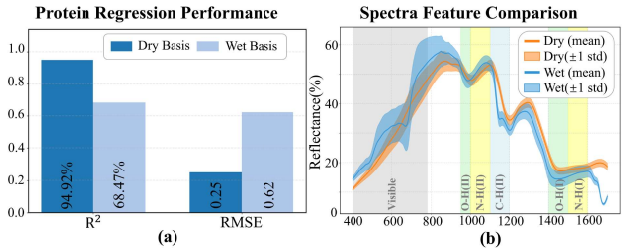

Figure 4: Effect of grain moisture on (a) protein prediction performance and (b) VNIR spectra of dry and wet samples.

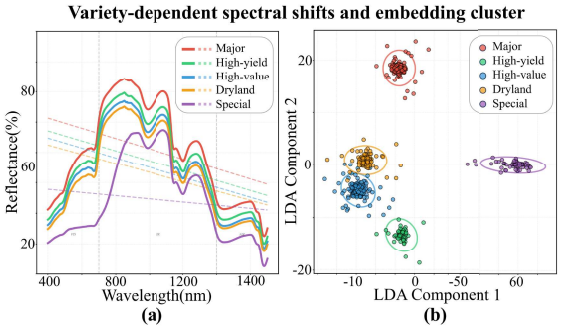

Figure 5: VNIR mean spectra (a) and low-dimensional embeddings (b) of five wheat variety groups derived from ten cultivars.

3.2.3 Generalization across varieties. Kernel hardness, pigment content, and genotype-driven microstructure shift baselines and band centers in VNIR [10, 11, 35]. Our ten wheat cultivars can be grouped into five major breeding types, which show systematic spectral differences. As illustrated in Figure 5, mean VNIR curves for each group exhibit distinct continuum levels and linear baselines; the fitted dashed lines (computed by a per-group linear regression baseline = polyfit($\lambda, R(\lambda), 1$)) reveal slope differences across types. These variety-dependent shifts in baseline and band contrast lead to clearly separated clusters in the spectral embedding space. Consequently, traditional regression methods like PLS or standard CNNs, which assume independent and identically distributed data, often fail to distinguish these domain-specific shifts from chemically relevant features, resulting in poor generalization on unseen cultivars. Cultivar-invariant yet chemically informative representations are therefore essential for scalable deployment.

3.2.4 Spectral variance from field scattering and geometry. Handheld, on-ear measurements face uncontrolled probe angles, stand-off distance, specular glints and shadowing, and in practice deliver lower analytical performance than laboratory NIRs [24]. Multi-scale scattering from glumes and overlapping kernels changes path lengths and band shapes; sparse bands amplify baseline drift. Hardware interventions such as more homogenized illumination reduce variance but cannot fully linearize geometry-induced nonlinearities.

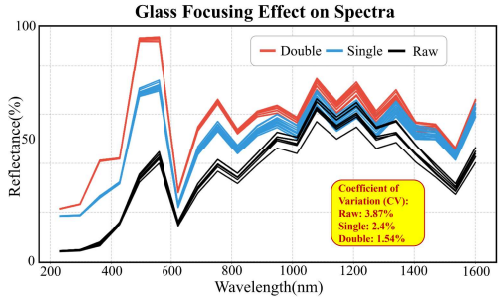

Figure 6: Effect of hardware-level optical scatter-normalization on field spectra. The plot compares the coefficient of variation (CV) across different optical setups: "Raw" (no diffuser), "Single" (single-sided diffuser), and "Double" (our proposed double-sided frosted glass). The double-layer design (red) effectively suppresses spectral variance from 3.87% (Raw) to 1.54%.

To address this, we propose a hardware-algorithm co-designed approach. At the hardware level, we introduce a double-sided frosted glass diffuser (the "optics") to physically homogenize light and suppress scattering variance at the source. As shown in Figure 6, this hardware intervention alone significantly reduces the spectral coefficient of variation (CV) compared to single-sided ("Single") or no diffuser ("Raw") setups. At the algorithmic level, this cleaner signal is fed into a multi-scale CNN, which is jointly optimized with reconstruction and task losses to further stabilize residual baseline drifts.

4 System Design

This section details the architecture of WheatScout, which is co-designed to provide robust estimates for these two key metrics directly in the field. The overall pipeline, illustrated in Figure 7, transforms raw, noisy sensor readings from our low-cost hardware into reliable compositional estimates by systematically addressing the core challenges outlined above.

The system's design philosophy centers on addressing three critical challenges through specialized components:

- **Hardware-Algorithm Collaborative Scattering Correction (§ 4.1).** Combining optical engineering with computational methods to mitigate scattering effects.
- **Moisture-First Dry-Basis Modeling (§ 4.2).** Addressing water interference through explicit moisture estimation and spectral normalization.
- **Domain-Invariant Feature Learning (§ 4.3).** Ensuring generalization across wheat varieties through adversarial training techniques.

4.1 Hardware-Algorithm Collaborative Scattering Correction

When performing handheld spectroscopic measurements on intact wheat ears in the field, minute changes in probe angle, pressure, and distance can significantly alter the path length of photons within the sample, leading to severe scattering effects. These effects manifest spectrally primarily as large baseline drift and signal distortion that are difficult to distinguish from the actual chemical composition, and are the main source of noise causing non-reproducible measurement results. To fundamentally address this issue, WheatScout employs a co-design strategy. Firstly, it suppresses scattering at the physical level through hardware optical design, and then adaptively compensates for residual, sample-related scattering using a dedicated algorithm.

4.1.1 Hardware-Level Scattering Suppression. The core design philosophy at the hardware level is to provide **a consistent and pristine optical environment for each measurement through a fixed geometric structure and active light control**, thereby minimizing the impact of external environmental and operational variations. This not only reduces the compensation burden on subsequent algorithms but also improves the quality of the original signal. Our innovations are mainly reflected in two key optical designs: fixed geometric structure and double-sided diffusion scheme.

Fixed Geometric Structure. To eliminate ambient light interference and lock in measurement conditions, we designed a closed probe housing with an internal baffle and a high-precision knife-edge aperture. This design ensures that the internal ambient light intensity is below 1 lux during measurement, thus isolating external stray light. Its innovation lies in using a kinematically constrained cylindrical base to strictly fix the relative geometry between the probe and the wheat ear at a normal incidence condition. This design eliminates the main scattering variations caused by changes in handheld distance and angle, which is impossible to achieve with open probes or simple contact probes.

Double-Sided Diffusion Scheme. Despite fixed geometric relationships, light still undergoes unpredictable specular reflection and scattering on the complex surfaces of wheat ears (such as husks and awns). To address this, we introduced a double-sided frosted glass window as a core innovative component (Fig. 8). Its working principle utilizes the microscopic roughness of the glass surface to perform Lambertian diffusion of light. Specifically, this glass window achieves two key functions: Homogenization of incident light: transforming the originally concentrated LED point light source into a spatially uniform illumination field. Diffusion of reflected light: performing secondary diffusion on the light reflected from the wheat ear surface, which already carries scattering information, smoothing out signal spikes caused by localized specular reflection and strong scattering. This process can be described by the Lambert cosine law, where the scattered light intensity $I(\theta)$ is proportional to the cosine of the observation angle θ, i.e., $I(\theta) = I_0 \cdot \frac{\rho}{\pi} \cdot \cos\theta$, where I_0 represents the incident light intensity and ρ represents the surface reflectivity. Ultimately, we obtain a more uniform measurement spot with a larger coverage area, reducing the influence of local scattering.

4.1.2 Algorithm-Level Scattering Compensation. After hardware optimization, the spectral signal still retains scattering effects caused by differences in the microstructure of the wheat ear itself, such as the waxy coating of the husk and the texture of the epidermis. The

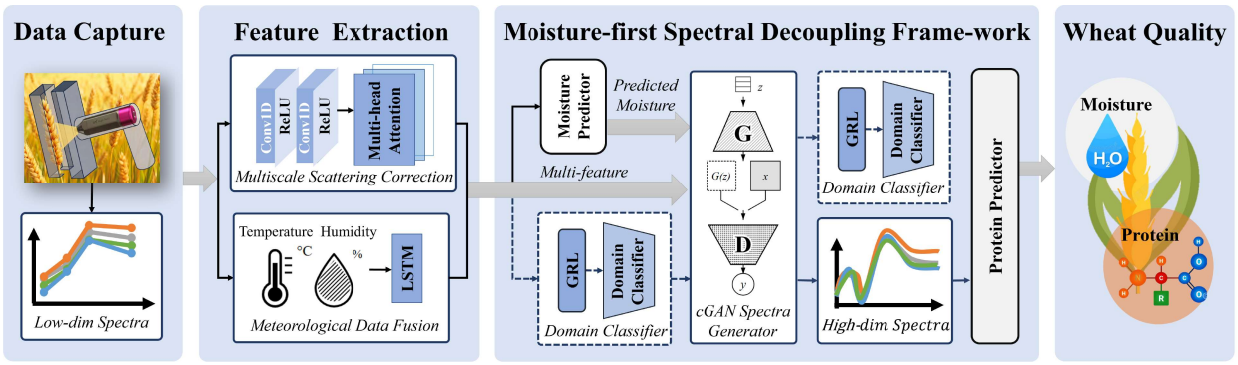

Figure 7: Overall architecture of WheatScout, including scattering correction, moisture-first dry-basis reconstruction, and domain-invariant protein/moisture prediction.

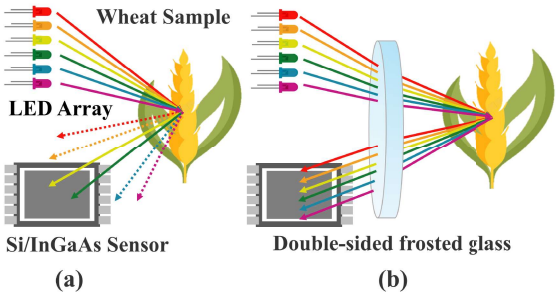

Figure 8: Optical principle of dual-sided frosted glass for scattering correction.

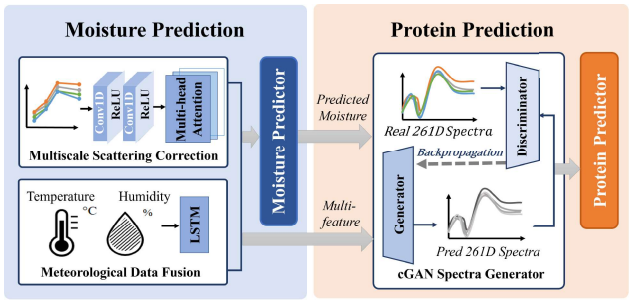

Figure 9: Overview of the cascaded dry-basis modeling framework.

task of the algorithm layer is to further learn feature representations that are insensitive to changes in residual scattering from this relatively clean signal. We employ a convolutional neural network (CNN) with multi-scale one-dimensional convolutional kernels (3×1 and 5×1) as the backbone for feature extraction. The design logic is that different scattering modes primarily affect the spectral baseline (low-frequency components), while the true absorption peaks exhibit local fluctuations (high-frequency components). The multi-scale convolutional kernels work collaboratively—the smaller kernel (3×1) focuses on extracting local absorption peak shape features and is insensitive to baseline drift; the larger kernel (5×1) can perceive a wider range of spectral shape changes, helping the model distinguish between true broad absorption bands and slow baseline changes caused by scattering. In this way, the CNN learns to focus on spectral details related to chemical composition while ignoring spectral distortions caused by scattering.

4.2 Moisture-First Dry-Basis Modeling

When directly measuring whole wheat ears in the field, the high moisture content, typically ranging from 10% to 40%, presents a significant challenge for accurate protein estimation. The strong absorption of moisture not only saturates critical wavelengths (e.g., around 1450 nm) but also distorts the entire spectral profile through nonlinear changes, effectively masking weak protein signals. Traditional approaches that attempt to regress protein content directly from the "wet" spectrum require the model to isolate weak signals from a heavily contaminated background, leading to substantial biases and poor generalization. The results in Figure 4(a) show that protein estimation on dry wheat samples is significantly more accurate than on fresh wheat samples. Inspired by this, we propose a "separate first, analyze later" strategy to decouple these intertwined factors. Instead of a direct end-to-end prediction, we model the problem as a two-stage process. First, we aim to recover latent "dry" spectral features from moist observations, thus effectively eliminating moisture interference computationally. Second, protein concentrations are predicted based on these recovered, moisture-normalized spectra.

4.2.1 Physics-Constrained Generative Reconstruction. To translate the proposed "separate-then-analyze" strategy into a deployable algorithm, we designed a conditional generation framework specifically tailored for spectral reconstruction.

First, we used a low-cost spectral sensor array capable of capturing 44 specific narrowband channels covering key absorption peaks of moisture and proteins. Furthermore, considering the thermodynamic sensitivity of moisture absorption, we introduced the environmental variable X_{weather} (temperature, humidity) as auxiliary covariates. These variables are used to calibrate the nonlinear spectral response drift caused by environmental conditions, providing the necessary physical background for reconstruction.

However, establishing a mapping from these sparse, distorted wet inputs to ideal dry spectra presents a significant modeling

challenge. A natural alternative is to directly regress the dry-basis spectrum from the wet spectrum using a deterministic model such as an MLP, a standard 1D-CNN, or an autoencoder trained only with pointwise reconstruction losses. However, in our setting such models are prone to an averaging effect: when multiple wet spectra with slightly different moisture states and scattering conditions correspond to similar dry-basis targets, minimizing only an L_1/L_2 loss encourages the network to produce over-smoothed outputs that match the average intensity but attenuate narrow, diagnostically important spectral structures. This issue is particularly harmful for protein estimation, because protein-related cues in fresh wheat ears are already much weaker than moisture-driven variations and often appear as subtle local deviations rather than dominant peaks.

We therefore adopt a conditional GAN to regularize the reconstruction toward the manifold of realistic dry-basis spectra. In our design, the adversarial loss complements the reconstruction loss by penalizing spectrally implausible outputs even when their pointwise error is small. This encourages the generator to preserve local peak shape, relative band contrast, and high-frequency structure that would otherwise be smoothed out by purely deterministic regression. In other words, the cGAN is not introduced for visual realism, but for spectral-shape fidelity: it helps recover dry-basis spectra that are both numerically accurate and chemically informative for downstream protein prediction.

Figure 9 shows the architecture of our proposed cascaded dry-basis modeling framework. Specifically, a feature extractor F_θ firstly learns a joint representation H from the original 44-dimensional wet spectrum X_{wet} and simultaneously acquired weather data X_{weather}. The weather data provides the necessary context for understanding the impact of temperature on moisture state. Next, the moisture prediction module M_ϕ estimates the moisture content $\hat{\alpha} = M_\phi(H)$ based on these joint features, which is the foundation for subsequent dry-based transformations.

The key innovation of the system lies in the conditional generator G_ψ, which fuses the joint feature H with the predicted moisture content $\hat{\alpha}$, performing a nonlinear mapping from wet to dry spectra: $\hat{S}_{\text{dry}} = G_\psi(H, \hat{\alpha})$. Unlike generic GANs, our architecture is tailored for spectral continuity. The generator G_ψ employs a U-Net-like structure based on 1D-CNNs with residual connections. This design allows the network to capture local spectral correlations (via convolution) while preserving high-frequency signal details (via residual links) during the upsampling from sparse 44-band inputs to dense 261-band dry spectra.

This generator transforms sparse 44-dimensional LED readings into dense 261-dimensional dry spectra. To ensure the physical plausibility of the generated spectra, we introduce a discriminator D_ξ, which forms an adversarial training framework with the generator. Its objective function is:

$$\min_G \max_D \mathbb{E}[\log D(S_{\text{dry}})] + \mathbb{E}[\log(1 - D(G_\psi(H, \hat{\alpha})))] \tag{2}$$

The rationale behind this adversarial mechanism is to impose a "manifold constraint." While the reconstruction loss (MSE) ensures numerical proximity, the discriminator forces the generator to produce spectra that strictly conform to the topological distribution of real dry wheat spectra. This effectively prevents the model from generating smooth but chemically meaningless curves, ensuring that the recovered protein-specific absorption features are physically valid.

The entire framework is collaboratively optimized through end-to-end multi-task learning. The total loss function is defined as:

$$\mathcal{L}_{\text{total}} = \lambda_1 \cdot \mathcal{L}_{\text{moisture}} + \lambda_2 \cdot \mathcal{L}_{\text{GAN}} + \lambda_3 \cdot \mathcal{L}_{\text{protein}} \tag{3}$$

In implementation, L_{GAN} consists of an adversarial term and a reconstruction term, i.e., $L_{GAN} = \lambda_{adv} L_{adv} + \lambda_{rec} L_{rec}$, where L_{rec} is computed as the mean squared error between generated and reference dry-basis spectra. This multi-task design facilitates knowledge sharing: accurate moisture estimation provides the necessary physical condition ($\hat{\alpha}$) for the GAN, while the high-quality dry spectral reconstruction simplifies the feature space for the final protein prediction.

4.3 Domain-Invariant Feature Learning Strategy

As described in Section 3.2.3, different wheat varieties exhibit distinct near-infrared spectral profiles due to morphological and microstructural variations, which manifest as systematic baseline drift and changes in absorption band contrast (Figure 5). Consequently, models trained on specific varieties may incorrectly learn to associate these morphology-related spectral shifts with protein content. When applied to new cultivars, the predictive performance of such models deteriorates significantly, necessitating recalibration or retraining with new, often scarce, labeled data [1, 30]. This data dependency poses a major barrier to scalable deployment.

To overcome this limitation and achieve robust generalization, the model must learn a feature representation that is invariant to these variety-induced spectral shifts while remaining sensitive to the underlying chemical information related to protein. This subsection details our strategy, which synergistically combines dual adversarial training and attention mechanisms, designed to explicitly separate morphology-related spectral variations from protein-related signals.

4.3.1 Formalization of Domain Shift Problem. When different wheat varieties (domains) $D_1, D_2, \ldots, D_K$ correspond to different data distributions $P(X, Y)$, a domain shift problem arises. In spectral analysis, this manifests as both covariate shift (input distribution changes) and conceptual shift (conditional distribution changes). The core challenge is to learn a representation that captures protein-related features while remaining invariant to variety-specific morphological variations. To address this, we propose a Synergistic Disentanglement Strategy that combines Dual Domain Adversarial Training with Attention-Enhanced Feature Extraction. These two components operate on a complementary principle: it simultaneously enforces domain invariance to suppress variety-specific noise and enhances spectral sensitivity to capture subtle biochemical signals.

4.3.2 Dual Domain Adversarial Training. Our system implements domain adversarial training at both the feature level and the generated spectrum level using a Gradient Reversal Layer (GRL). The domain adversarial loss at the feature level is:

$$\mathcal{L}_{\text{domain1}} = -\mathbb{E}[\log D_1(F_\theta(X))] \tag{4}$$

And at the spectral level:

$$\mathcal{L}_{\text{domain2}} = -\mathbb{E}[\log D_2(G_\psi(H, \hat{\alpha}))] \tag{5}$$

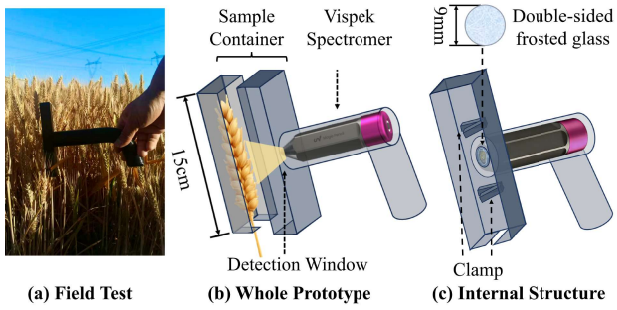

Figure 10: Prototype of WheatScout, which consists of a commercial Vispek spectrometer and an enclosure.

The rationale for this dual-level design is to enforce a strict decoupling of morphology and chemistry. The feature-level adversary prevents the encoder from retaining variety-identifying information (like awn patterns). The spectral-level adversary ensures that the generated dry spectrum represents a universal wheat ear rather than a specific variety, preventing the GAN from hallucinating variety-specific artifacts.

4.3.3 Attention-Enhanced Feature Extraction. While adversarial training removes unwanted information, it risks removing too much. To counterbalance this, we employ a multi-head self-attention mechanism:

$$\text{Attention}(Q, K, V) = \text{softmax}\left(\frac{QK^T}{\sqrt{d_k}}\right)V \tag{6}$$

Unlike standard CNNs which process local neighborhoods, the self-attention mechanism models global dependencies across the entire spectral range. This is crucial because chemical bonds (like N-H and C-H) often exhibit coupled absorption peaks at distant wavelengths (e.g., overtones and combination bands). The attention mechanism allows the model to lock on to these stable, cross-variety chemical fingerprints, ensuring that the features used for prediction are chemically relevant even after the adversarial purging of morphological information.

By integrating these mechanisms, the system effectively filters out the morphological noise introduced by variety differences while preserving and enhancing the weak spectral signals of protein, achieving robust generalization.

5 Implementation

Our WheatScout system integrates a commercial spectrometer (26×99.5 mm, 65 g) within a custom, light-absorbing 3D-printed enclosure (Figure 10). The device is controlled via a smartphone app and features a 500 ms acquisition cycle.

Enclosure Design and Rationale. Rather than serving merely as a passive protective shell, the enclosure is designed as an active optical-mechanical component to stabilize the illumination geometry and reduce acquisition variance at the source. To achieve both low-cost field manufacturability and optical isolation, the housing is 3D-printed using matte black PLA, which inherently minimizes internal stray reflections.

The enclosure features a shallow rectangular sample chamber sized specifically to accommodate intact wheat spikes. The lower shell has an outer footprint of approximately 152×32 mm, while the effective measurement cavity is 145×25 mm with an inner height of 14.2 mm. Within this cavity, two plastic clips are glued directly adjacent to both sides of the optical window. This arrangement firmly secures the wheat spike, keeping it almost flush against the optical interface. This mechanism maintains a consistent, minimal stand-off distance, preventing tilt, sagging, or off-axis motion during handheld operation.

To ensure a fixed normal-incidence geometry, the spectrometer is kinematically constrained within a coaxial cylindrical probe sleeve. The sleeve has an inner diameter of 27.7 mm, an outer diameter of 32.4 mm, and a straight insertion length of 80 mm. This passive guide securely constrains the probe's orientation and insertion depth, eliminating lateral drift and improving repeatability across different users. At the center of the chamber, the optical path is defined by a 12-mm circular aperture passing through a 2.5-mm-thick section, yielding an effective 9-mm probe spot centered along the spike axis to average signals across kernels, glumes, and awns.

To further enhance signal fidelity, a double-sided frosted glass diffuser is installed at the probe aperture. This element homogenizes the illumination field and effectively suppresses scattering artifacts and specular reflections from the waxy kernel cuticles. When closed, the assembly forms an optically sealed environment, maintaining ambient light infiltration below 1 lux. Ultimately, these hardware design choices—the fixed probe geometry, the homogenizing diffuser, and the light-tight matte chamber—suppress geometry-induced baseline drift and scattering variance before the spectra reach the regression model, significantly simplifying the downstream machine learning task.

Spectrometer Selection and Parameters. We selected the Vispek Magic Pencil for its broad 200–1700 nm coverage. The device utilizes 22 LED bands spanning UV–VIS–NIR, each captured with dual-drive (high/low) intensity, yielding a 44-dimensional output for robust spectral reconstruction. This extensive range covers critical absorption bands for wheat grain protein and moisture quantification.

The instrument offers 33 nm optical resolution with high responsivity (approx. 0.96 A/W at 940 nm and 0.95 A/W at 1550 nm). This performance ensures a high Signal-to-Noise Ratio (SNR) across the target spectral windows, facilitating stable calibration for in-field applications.

6 Performance Evaluation

6.1 Study Setup

6.1.1 Metrics. We reformulate the pre-harvest wheat quality assessment task into a supervised regression task that maps in-field handheld spectra to laboratory measurements of moisture and protein. To explicitly evaluate generalization to unseen varieties, we adopt a variety-based cross-validation scheme on the 15-day, 10-variety dataset described above. In each fold, two wheat varieties are held out for testing, while the remaining eight varieties are used for training. This setup probes whether the model has learned a variety-agnostic mapping from spectra to moisture and protein, rather than memorizing variety-specific signatures. Predictive performance

Table 2: Stratified 5-fold cross-validation performance comparison between WheatScout and baselines.

Trait	Models	R^2	RMSE
Protein	44D+mPLS	58.97%	0.71
	44D+ANN	65.53%	0.65
	Recon+mPLS	74.38%	0.56
	Recon+ANN	78.67%	0.51
	WheatScout	92.38%	0.31
	261D+mPLS	94.32%	0.26
	261D+ANN	94.92%	0.25
Moisture	44D+mPLS	82.18%	5.82
	44D+ANN	85.75%	5.21
	WheatScout	97.35%	2.25

is quantified using the coefficient of determination (R^2) and root mean squared error (RMSE), reported separately for moisture and protein. All metrics are computed exclusively on the held-out varieties and then averaged across folds to obtain robust estimates of cross-variety field performance.

6.1.2 Baselines. We compare our method against three groups of baselines that reflect both current commercial practice and state-of-the-art spectral modeling.

- **Handheld Baselines (44D + mPLS / ANN).** We predict moisture and protein from 44D spectra using modified Partial Least Squares (mPLS) and Artificial Neural Networks (ANN), matching standard options in commercial NIR calibration software like FossCalibrator [12]. As no prior learning algorithms exist specifically for this low-dimensional, LED-based handheld device for pre-harvest wheat, we adopt these commercial models as industry-relevant baselines.
- **Reconstruction Baselines (Reconstructed 44D + mPLS / ANN).** To evaluate our spectral reconstruction algorithm, we use a SOTA method [15] to map 44D spectra to 261D, then apply mPLS and ANN to predict protein from the reconstructed spectra. Results are reported only for protein, as our moisture regression models were not trained on 261D data.
- **Lab Baselines (261D + mPLS / ANN).** We predict protein using mPLS and ANN on 261D dry matter spectra. Moisture results are omitted as the 261D data was not used for moisture training. This setup represents an ideal high-resolution laboratory scenario, highlighting the performance gap between benchtop instruments and our low-dimensional handheld sensor.

6.2 Overall Performance

We first evaluate WheatScout's performance in predicting moisture and protein by conducting 5-fold cross-validation stratified by wheat variety.

6.2.1 Baselines Comparison. Table 2 compares the WheatScout system against baseline methods for protein and moisture prediction. The low-dimensional 44D+ANN model achieved results of 65.53% and 85.75%. Incorporating spectra reconstructed via the SOTA model [15] (Recon+ANN) increased the protein prediction

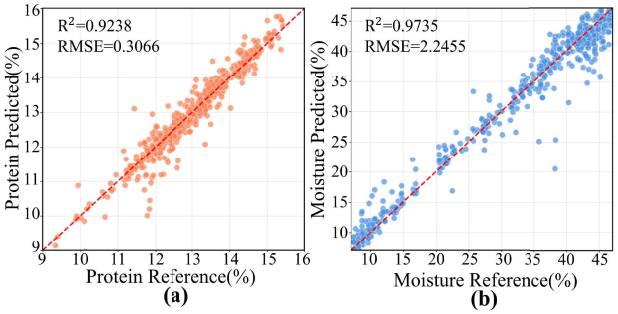

Figure 11: Scatter plots comparing actual versus predicted values. (a) The prediction performance for Protein. (b) The prediction performance for Moisture. The red dashed line represents the ideal prediction.

R^2 to 78.67%. The WheatScout method achieved an R^2 of 92.38% for protein and 97.35% for moisture, results comparable to those obtained from high-dimensional data.

6.2.2 WheatScout Performance. Figure 11 displays the regression results for protein and moisture. In Figure 11 (a), the predicted protein values align with the ground truth along the identity line ($y = x$). Similarly, Figure 11 (b) shows that moisture predictions follow the ideal trajectory across the value range. These plots indicate a linear correspondence between the predicted outputs and the actual measurements.

6.2.3 Effect of Excluding Visible Bands. One concern in crop sensing is that visible-range differences, such as ear color or cultivar-specific pigmentation, may provide shortcut cues that are not causally related to protein or moisture. To explicitly test if WheatScout is biased toward external color, we evaluated an NIR-only model (≥ 750 nm) to completely remove visible cues. As Table 3 shows, this model maintained high accuracy (Protein R^2: 89.54%, Moisture R^2: 95.82%). This provides explicit experimental evidence that predictions rely on internal chemical bonds (e.g., O-H, N-H) rather than merely correlating with superficial color. The slight performance gain from including visible bands suggests they provide complementary structural or scattering information, rather than acting as the primary predictive signal.

Table 3: Performance comparison of WheatScout using the full spectrum versus NIR-only bands to evaluate the influence of external color.

Trait	Bands Used	*R2*	RMSE
Protein	Full Spectrum (WheatScout)	92.38%	0.31
	NIR Only (≥ 750 nm)	89.54%	0.36
Moisture	Full Spectrum (WheatScout)	97.35%	2.25
	NIR Only (≥ 750 nm)	95.82%	2.85

Table 4: Performance comparison of WheatScout with its ablation variants and SOTA methods. The best results are highlighted in bold.

Model	Protein		Moisture	
	R^2	RMSE	R^2	RMSE
SOTA	78.67%	0.51	85.75%	5.21
w/o Dry-Basis	82.15%	0.47	94.12%	3.34
w/o Domain-Adv	88.40%	0.38	96.19%	2.69
w/o Attention	90.12%	0.35	96.81%	2.46
WheatScout (Ours)	**92.38%**	**0.31**	**97.35%**	**2.25**

6.3 Ablation Study

To evaluate the contribution of the core components of WheatScout, we conducted an ablation study examining the impact of moisture-first dry-basis modeling (cGAN), domain-invariant feature learning (domain adversarial training), and the attention mechanism. The variants compared are as follows:

- **w/o Dry-Basis**: Removes the cGAN generator; predicts protein directly based on wet-basis spectra without moisture decoupling.
- **w/o Domain-Adv**: Removes domain adversarial training; optimizes solely via regression loss without variety-invariance constraints.
- **w/o Attention**: Removes the self-attention module; relies only on the CNN backbone for feature extraction.

The quantitative results are summarized in Table 4.

6.3.1 Effectiveness of Moisture-First Dry-Basis Modeling. Removing the dry-basis module (**w/o Dry-Basis**) caused the most significant performance drop, with protein prediction RMSE rising to 0.47 and R^2 decreasing by over 10%. This confirms that high moisture content in fresh wheat ears generates strong nonlinear interference, masking weak protein spectral features. Lacking the standardized dry-basis reference generated by cGAN, the model struggles to decouple protein signals from dominant water absorption bands, leading to estimation bias. The dry-basis design is fundamental for high-precision in-field detection.

6.3.2 Effectiveness of Domain-Invariant Learning. For **w/o Domain-Adv**, the protein R^2 dropped to 88.40%. This indicates that without dual domain adversarial training, the model is prone to overfitting morphological characteristics of specific varieties, such as glume texture, color or awns. The domain adversarial mechanism forces the feature extractor to discard these variety-specific shortcut features and focus on stable chemical bond spectral patterns, thereby ensuring robustness and generalization across different wheat varieties.

6.3.3 Effectiveness of Attention Mechanism. Removing the attention module (**w/o Attention**) resulted in a slight performance decline. While the multi-scale CNN backbone captures local spectral shapes, the attention mechanism dynamically weights band importance via global context. This module suppresses irrelevant background noise and highlights key absorption regions, such as N-H and C-H overtones, further refining features to enhance final detection accuracy.

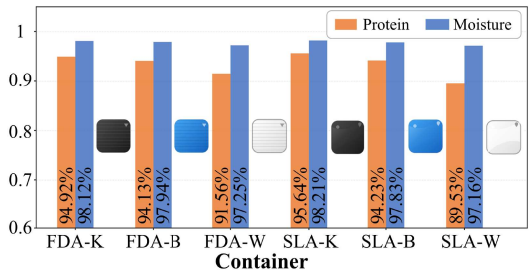

Figure 12: R^2 scores with varying container color and material.

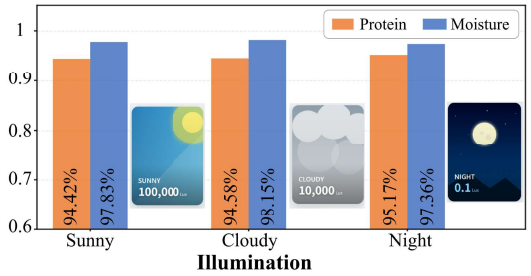

Figure 13: R^2 with varying ambient illumination.

6.4 Impacts on Experimental Conditions

We evaluated the robustness of WheatScout under various conditions, with 15 samples per case. Models were trained using in-field data. Unless noted otherwise, the representative wheat variety Jimai 22 was used throughout the study.

6.4.1 Varying container color and material. We tested different 3D-printed sample holders, including FDM and SLA methods and three colors (Figure 12). Darker containers maintained high R^2 scores, while white smooth surfaces caused a slight protein R^2 drop to 0.89, likely due to light scattering.

6.4.2 Varying Ambient Illumination. The system was evaluated under illumination levels of sunny, cloudy and night. As shown in Figure 13, performance was stable under all lighting conditions. This is due to the fully enclosed design, which isolates the measurements from ambient light.

6.4.3 Varying Ambient Temperature. We evaluated temperature robustness in a climate chamber at 60% RH across a temperature range from 10°C to 50°C(Figure 14). With stable performance (protein and moisture R^2 both remaining above 90%) from 10 to 40°C, the system fully encompasses the critical harvest windows for spring (15-25°C) and winter wheat (20-30°C). Significant degradation occurred only at 50°C, likely due to sensor thermal noise or moisture evaporation in extreme heat.

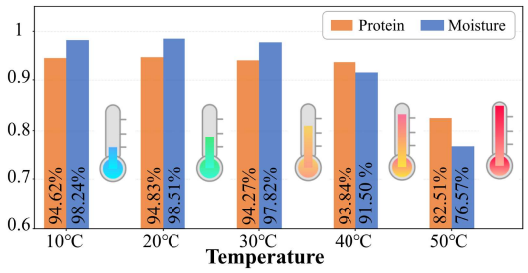

Figure 14: R^2 with varying ambient temperature.

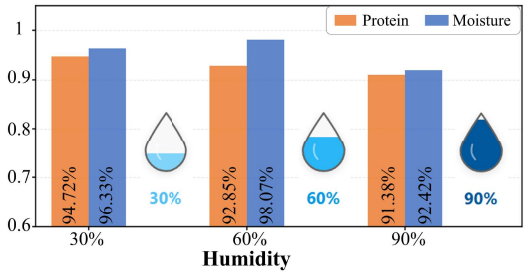

Figure 15: R^2 with varying ambient humidity.

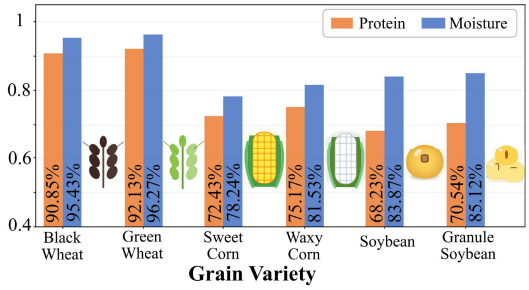

Figure 16: R^2 with varying grain variety.

6.4.4 Varying Ambient Humidity. We evaluated humidity robustness in a climate chamber at a constant 25°C across gradients of 30%, 60%, and 90% RH. As shown in Figure 15, performance decline was observed at 90% RH, likely due to condensation interference.

6.4.5 Varying Grain Variety. We extended our evaluation to include highly distinctive Green and Black Wheat, as well as Sweet Corn, Waxy Corn, Soybean, and Small Soybean (Figure 16). Despite the unique characteristics of the special wheat varieties, the model maintained high accuracy. However, predictions for non-wheat crops showed an anticipated decrease in performance (Protein R^2: 68%–75%), reflecting the biological spectral differences between wheat and coarse grains.

6.4.6 Varying User Operation. To evaluate the device's performance in authentic usage scenarios, we recruited five local farmers from Tai'an, Shandong, China. Following a brief training session,

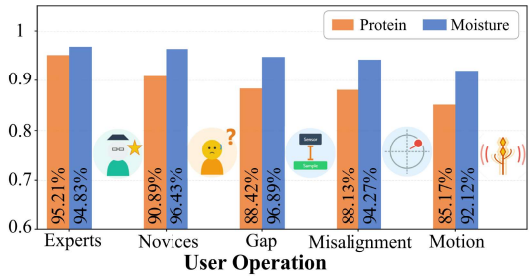

Figure 17: R^2 with varying user operation.

each participant collected three samples without supervision. Additionally, to rigorously assess system robustness, we manually simulated three potential interference conditions: a 5 mm vertical gap (z-axis), a 5 mm horizontal offset (x-y axis), and slight shaking. As shown in Figure 17, While the participants demonstrated competent device operation, the simulated interference scenarios resulted in a noticeable decline in detection accuracy. Tilting the probe changes the optical path length and the angle of incidence, which can introduce specular reflection components that distort the spectral baseline. Furthermore, external motion or ambient light fluctuations introduce stochastic noise into the spectral data.

7 Related Work

In this section, we briefly review existing works related to wheat quality sensing, spectral learning under moisture, and low-cost spectral systems.

Wheat quality sensing. Benchtop NIR and hyperspectral instruments are the gold standard for quantifying grains [4, 9, 34]. However, these stationary systems are not suitable for on-site testing. Handheld devices bring NIR to the farm but remain expensive and typically require threshed grain [1]. At larger scales, UAV and ground-based imaging estimate canopy-level traits [14, 29, 31], but these methods rely on indirect vegetation indices and lack the resolution for direct physicochemical analysis. WheatScout bridges this gap by enabling pre-harvest, on-ear estimation using a handheld device designed for smallholders.

Spectral learning and data diversity. Conventional chemometrics (e.g., PLS) and preprocessing methods (e.g., SNV, MSC) are effective for seeds at stable storage moisture [30]. In pre-harvest settings, however, strong water absorption bands obscure protein signals, and linear models struggle to disentangle these variations. While some studies attempt to broaden calibration ranges [22, 23], they are often validated on narrow datasets that fail to capture real-world variability. In contrast, our study is grounded in a comprehensive dataset. We cover 10 distinct wheat varieties representing the genetic diversity of Northern China, including 5 important types, spanning a range of moisture (7.36%–46.85%) and protein contents (9.27%–15.44%). This extensive data scale ensures our approach is validated against realistic genetic variations and physiological states, overcoming the limitations of prior small-scale studies.

Low-cost spectral systems. Compact VIS-NIR sensors based on LEDs have reduced hardware costs for grain analysis [20], yet they generally target single analytes in threshed kernels. Mobile

computing systems using smartphone cameras or diffraction gratings have shown promise for homogeneous samples [15, 16, 28] but rarely address the complex scattering and geometric artifacts of intact cereal ears. WheatScout advances this domain by co-designing diffuser-based optics with a robust reconstruction pipeline, achieving on-ear quality estimation at a hardware cost two orders of magnitude lower than laboratory instruments.

8 Discussion

In this section, we discuss the current limitations of WheatScout and outline potential pathways for future refinement, focusing on real-world robustness, application expansion, and scalability.

Practical Robustness and Hardware Evolution. Although WheatScout demonstrates stable performance under controlled field trials, long-term real-world deployment introduces new challenges. Mechanically, the enclosure's sealing quality may degrade over time due to repeated opening, dust accumulation, or wear, potentially increasing ambient light leakage. Furthermore, user operation introduces residual variability, such as incomplete chamber closure, off-center spike placement, or excessive motion. Environmentally, extreme temperature and humidity can alter LED emission spectra, sensor noise characteristics, and the surface moisture of the wheat ear.

To address these practical issues, future hardware iterations will incorporate compressible light-blocking gaskets, magnetic closures, and replaceable diffuser cartridges to maintain optical consistency. At the system level, we plan to integrate auxiliary low-cost sensors (e.g., temperature, humidity, and simple RGB cameras). This multi-modal sensor fusion will not only provide context-aware compensation for environmental drift but also help overcome the inherent resolution limits of sparse multi-LED sensors when absorption bands overlap. Finally, at the algorithmic level, implementing signal-quality monitoring and confidence-aware predictions will allow the system to detect unreliable acquisitions and prompt users to re-measure.

Scope, Trait, and Crop Expansion. While WheatScout achieves strong domain adaptation across various cultivars, our current evaluation primarily focuses on protein and moisture in Chinese wheat, with limited representation in the very low-protein range ($< 10\%$). Validating the system across a broader spectrum of international cultivars is essential for global deployment. Beyond protein and moisture, future work will explore inferring traits indirectly linked to spectral signatures, such as grain hardness, sedimentation value, and test weight, by leveraging the latent biochemical correlations captured by our representation learning module.

Additionally, WheatScout currently lacks zero-shot generalization to structurally distinct grains like maize or soybean. Given the prevalence of crop rotation in smallholder farming, a single device capable of monitoring multiple crop types is economically vital. Future research will focus on developing a universal backbone using meta-learning strategies, enabling the system to rapidly adapt to the unique scattering profiles and chemical compositions of legumes and coarse cereals.

Cost Analysis for Large-Scale Field Deployment. A core motivation behind WheatScout is to democratize precision agriculture for individual smallholder farmers. In stark contrast to prohibitively expensive commercial spectrometers, the total hardware cost of the current WheatScout prototype is under $100, as it is built entirely from off-the-shelf multi-LED sensors, basic microcontrollers, and 3D-printed components.

For large-scale deployment across individual farms, transitioning to injection molding and custom integrated PCBs is projected to drive the unit hardware cost down to approximately $20–$30. Beyond the initial hardware investment, the operational and maintenance costs are strictly minimized. The device relies on low-power components suitable for standard rechargeable batteries, and the heavy computational load (e.g., the cGAN and representation learning models) is entirely offloaded to the cloud. This cloud-based architecture means farmers only need a basic entry-level smartphone for data transmission and visualization, eliminating the need for expensive on-device processors. Furthermore, modular designs, such as the replaceable diffuser cartridge, ensure that long-term maintenance remains affordable. Ultimately, this exceptionally low total cost of ownership makes it economically viable to deploy WheatScout as a distributed, high-density sensor network, empowering individual farmers with real-time, field-level crop analytics.

9 Concluding Remarks

We present WheatScout, a low-cost handheld system designed to empower smallholder farmers with precise pre-harvest wheat quality monitoring. By enabling real-time, on-ear estimation of protein and moisture, WheatScout assists growers in optimizing harvest timing to maximize crop value. To ensure robust performance in the wild, we collected in-situ spectral data from 600 living wheat ears over a 15-day pre-harvest period. The system employs a hardware-algorithm co-design to robustly decouple protein signals from strong moisture interference across diverse cultivars. Extensive field experiments demonstrate that WheatScout achieves high accuracy ($R^2 > 97\%$ for moisture and $R^2 > 92\%$ for protein) using sparse LED measurements. Ultimately, WheatScout democratizes data-driven decision-making, offering an accessible solution for precise quality management in non-industrial farming contexts.

Acknowledgments

This research is supported in part by RGC under Contract CERG 16204523, 16205824, AoE/E-601/22-R, SRFS2425-6S05 and Contract R8015.

We thank Dr. Qin Lu from the Core Facility of State Key Laboratory of Hybrid Rice of Wuhan University for her assistance with near-infrared grain analyzer.

References

[1] [n. d.]. GrainSense Go Handheld Grain Analyzer. https://grainsense.com/. GrainSense Ltd., accessed: 27 November 2025.

[2] G. Y. Abawi. 1993. A Simulation Model of Wheat Harvesting and Drying in Northern Australia. *Journal of Agricultural Engineering Research* 54, 2 (Feb. 1993), 141–158.

[3] Douglas S Alt, Pierce A Paul, Alexander J Lindsey, and Laura E Lindsey. 2019. Early Wheat harvest influenced grain quality and profit but not yield. *Crop, Forage & Turfgrass Management* 5, 1 (2019), 1–6.

[4] Amanda Teixeira Badaró, João Pedro Hebling e Tavares, Jose Blasco, Nuria Aleixos-Borrás, and Douglas Fernandes Barbin. 2022. Near infrared techniques applied to analysis of wheat-based products: Recent advances and future trends. *Food Control* 140 (2022), 109115. doi:10.1016/j.foodcont.2022.109115

[5] Keerthi Chadalavada, Krithika Anbazhagan, Adama Ndour, Sunita Choudhary, William Palmer, Jamie R. Flynn, Srikanth Mallleye, Sharada Pothu, Kodukula

Venkata Subrahamanya Vara Prasad, Padmakumar Varijakshapanikar, Chris S. Jones, and Jana Kholová. 2021. NIR instruments and prediction methods for rapid access to grain protein content in multiple cereals. *Plants* 10, 9 (2021), 1796.

[6] China Chamber of Commerce of I/E of Foodstuffs, Native Produce and Animal By-products (CFNA). 2025. [Trade Overview] Overview of China's Agricultural Export Trade (Jan-Dec 2024). https://www.cccfna.org.cn/maoyitongji/guobiemaoyi/ff80808194628b47019545490dea0e39.html Original source in Chinese.

[7] Zhenling Cui, Hongyan Zhang, Xinping Chen, Chaochun Zhang, Wenqi Ma, Chengdong Huang, Weifeng Zhang, Guohua Mi, Yuxin Miao, Xiaolin Li, et al. 2018. Pursuing sustainable productivity with millions of smallholder farmers. *Nature* 555, 7696 (2018), 363–366.

[8] C Daniel and E Triboı. 2002. Changes in wheat protein aggregation during grain development: effects of temperatures and water stress. *European Journal of Agronomy* 16, 1 (2002), 1–12.

[9] Zhenjiao Du, Wenfei Tian, Michael Tilley, Donghai Wang, Guorong Zhang, and Yonghui Li. 2022. Quantitative assessment of wheat quality using near-infrared spectroscopy: A comprehensive review. *Comprehensive Reviews in Food Science and Food Safety* 21, 3 (2022), 2956–3009. doi:10.1111/1541-4337.12958

[10] Václav Dvořáček, Lenka Štěrbová, Eva Matějová, Jana Bradová, and Jiří Hermuth. 2018. Reflectance Spectrometry as a Screening Tool for Prediction of Lutein Content in Diverse Wheat Species (*Triticum* spp.). *Food Analytical Methods* 11 (2018), 2579–2589.

[11] Chyngyz Erkinbaev, Kieran Derksen, and Jitendra Paliwal. 2019. Single kernel wheat hardness estimation using near infrared hyperspectral imaging. *Infrared Physics & Technology* 98 (2019), 250–255.

[12] FOSS 2020. *FossCalibrator™: Unleash the business potential of calibration development.* FOSS, Hillerød, Denmark. https://www.fossanalytics.com Product one-pager.

[13] Charles S. Gaines and William R. Windham. 1998. Effect of Wheat Moisture Content on Meal Apparent Particle Size and Hardness Scores Determined by Near-Infrared Reflectance Spectroscopy. *Cereal Chemistry* 75, 3 (1998), 386–391.

[14] Fenner H. Holman, Andrew B. Riche, Adam Michalski, March Castle, Martin J. Wooster, and Malcolm J. Hawkesford. 2016. High Throughput Field Phenotyping of Wheat Plant Height and Growth Rate in Field Plot Trials Using UAV Based Remote Sensing. *Remote Sensing* 8, 12 (2016), 1031. doi:10.3390/rs8121031

[15] Haiyan Hu, Qianyi Huang, and Qian Zhang. 2023. Babynutri: a cost-effective baby food macronutrients analyzer based on spectral reconstruction. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 7, 1 (2023), 1–30.

[16] Haiyan Hu, Yinan Zhu, Baichen Yang, Hua Kang, Shanwen Chen, and Qian Zhang. 2024. MeatSpec: Enabling Ubiquitous Meat Fraud Inspection through Consumer-Level Spectral Imaging. In *Proceedings of the 30th Annual International Conference on Mobile Computing and Networking*. 861–874.

[17] Poonam Jasrotia, Beant Singh, and Mohini Nagpal. 2022. Biology and management strategies of major insect-pests of wheat. In *New horizons in wheat and barley research: crop protection and resource management*. Springer, 283–307.

[18] P. I. Kalandarov. 2022. High-frequency moisture meter for measuring the moisture content of grain and grain products. *Measurement Techniques* 65 (2022), 297–303.

[19] Neeraj Kumar, Ganesh Upadhyay, Krishna Bahadur Chhetri, BR Harsha, Gulshan Kumar Malik, Ravindra Kumar, Poonam Jasrotia, Shiv Ram Samota, Nitesh Kumar, RS Chhokar, et al. 2023. Pre-and post-harvest management of wheat for improving the productivity, quality, and resource utilization efficiency. In *Wheat science*. CRC Press, 57–106.

[20] Yiming Liu, Donghang Li, Huaiming Li, Xiaoping Jiang, Yan Zhu, Weixing Cao, and Jun Ni. 2022. Design of a Phenotypic Sensor About Protein and Moisture in Wheat Grain. *Frontiers in Plant Science* 13 (2022), 881560. doi:10.3389/fpls.2022.881560

[21] Mingming Ma, Yingchun Li, Cheng Xue, Wei Xiong, Zhengping Peng, Xue Han, Hui Ju, and Yong He. 2021. Current Situation and Key Parameters for Improving Wheat Quality in China. *Frontiers in Plant Science* 12 (2021), 638525. doi:10.3389/fpls.2021.638525

[22] S. Mahesh, D. S. Jayas, J. Paliwal, and N. D. G. White. 2011. Identification of wheat classes at different moisture levels using near-infrared hyperspectral images of bulk samples. *Sensing and Instrumentation for Food Quality and Safety* 5 (2011), 1–9. doi:10.1007/s11694-010-9104-2

[23] Kamaranga H. S. Peiris, Scott R. Bean, Anuj Chiluwal, Ramasamy Perumal, and S. V. Krishna Jagadish. 2019. Moisture Effects on Robustness of Sorghum Grain Protein Near-Infrared Spectroscopy Calibration. *Cereal Chemistry* 96, 4 (2019), 678–688. doi:10.1002/cche.10164

[24] Kamaranga H. S. Peiris, Scott R. Bean, Xiaorong Wu, Sarah A. Sexton-Bowser, and Tesfaye Tesso. 2023. Performance of a Handheld MicroNIR Instrument for Determining Protein Levels in Sorghum Grain Samples. *Foods* 12, 16 (2023), 3101. doi:10.3390/foods12163101

[25] S Pepler, MJ Gooding, and RH Ellis. 2006. Modelling simultaneously water content and dry matter dynamics of wheat grains. *Field Crops Research* 95, 1 (2006), 49–63.

[26] Vincent Ricciardi, Navin Ramankutty, Zia Mehrabi, Larissa Jarvis, and Brenton Chookolingo. 2018. How much of the world's food do smallholders produce? *Global food security* 17 (2018), 64–72.

[27] CL Rosser, Paweł Górka, AD Beattie, HC Block, JJ McKinnon, HA Lardner, and GB Penner. 2013. Effect of maturity at harvest on yield, chemical composition, and in situ degradability for annual cereals used for swath grazing. *Journal of Animal Science* 91, 8 (2013), 3815–3826.

[28] Neha Sharma, Muhammad Shahzaib Waseem, Shahrzad Mirzaei, and Mohamed Hefeeda. 2023. MobiSpectral: Hyperspectral imaging on mobile devices. In *Proceedings of the 29th Annual International Conference on Mobile Computing and Networking*. 1–15.

[29] Olga S. Walsh, Juliet M. Marshall, Eva Nambi, Chad A. Jackson, Emmanuella Owusu Ansah, Ritika Lamichhane, Jordan McClintic-Chess, and Francisco Bautista. 2023. Wheat yield and protein estimation with handheld and unmanned aerial vehicle-mounted sensors. *Agronomy* 13, 1 (2023), 207. doi:10.3390/agronomy13010207

[30] Phil C. Williams. 2020. Application of chemometrics to prediction of some wheat quality factors by near-infrared spectroscopy. *Cereal Chemistry* 97, 5 (2020), 958–966. doi:10.1002/cche.10318

[31] Shurong Yang, Lei Li, Shuaipeng Fei, Mengjiao Yang, Zhiqiang Tao, Yaxiong Meng, and Yonggui Xiao. 2024. Wheat Yield Prediction Using Machine Learning Method Based on UAV Remote Sensing Data. *Drones* 8, 7 (2024), 284. doi:10.3390/drones8070284

[32] Dandan Ye, Laijun Sun, Borui Zou, Qian Zhang, Wenyi Tan, and Wenkai Che. 2018. Non-destructive prediction of protein content in wheat using NIRS. *Spectrochimica Acta Part A: Molecular and Biomolecular Spectroscopy* 189 (2018), 463–472.

[33] Shuhua Zhang, Yongliang Chen, Pengfei Zhao, Hongyan Zhang, and Zhenling Cui. 2020. Increasing wheat yield for smallholder farmers with technology transfer. *Agronomy Journal* 112, 6 (2020), 5096–5104.

[34] Shun Zhang, Shuliang Liu, Li Shen, Shujuan Chen, Li He, and Aiping Liu. 2022. Application of near-infrared spectroscopy for the nondestructive analysis of wheat flour: A review. *Current Research in Food Science* 5 (2022), 1305–1312. doi:10.1016/j.crfs.2022.08.019

[35] Jochen U. Ziegler, Martin Leitenberger, C. Friedrich H. Longin, Tobias Würschum, Reinhold Carle, and Ralf M. Schweiggert. 2016. Near-infrared reflectance spectroscopy for the rapid discrimination of kernels and flours of different wheat species. *Journal of Food Composition and Analysis* 51 (2016), 30–36.