FlourSpec: Cost-Effective Spectral Analysis for Trace-Level Detection of Flour Adulteration

Shanwen Chen[1], Haiyan Hu[1], Yinan Zhu[1], Qian Zhang[1*]
[1]Department of Computer Science and Engineering, The Hong Kong University of Science and Technology
schenfh@cse.ust.hk, hhuap@cse.ust.hk, yzhudf@connect.ust.hk, qianzh@cse.ust.hk

ABSTRACT

Flour adulteration poses significant health risks and economic losses for consumers, yet current detection methods are often impeded by high costs, limited sensitivity, and impractical laboratory requirements, leaving end users vulnerable. The difficulty is exacerbated by the need for extremely low detection limits, often at parts per million (ppm) levels, to identify adulterants. To address these challenges, we introduce FlourSpec, a low-cost and user-friendly system designed for on-site detection of flour adulteration at the ppm level. FlourSpec employs a spectral reconstruction algorithm to extract pertinent information from coarse-grained spectral data collected using a low-cost spectrometer. Recognizing that existing spectral reconstruction algorithms often struggle with significant errors that obscure trace adulterant information, we develop a novel end-to-end architecture that balances spectral fidelity and classification performance while minimizing reconstruction error propagation. Additionally, we incorporate a hybrid attention mechanism to capture harmonic correlations within the spectra, effectively suppressing cross-band reconstruction errors. A supervised contrastive learning module is also incorporated to enhance the discriminability of trace features. Experimental evaluations demonstrate that FlourSpec achieves 97.44% accuracy in detecting multiple adulteration types and 93.12% accuracy in identifying various BPO concentrations. Notably, it is only 0.63% lower than expensive solutions and 18.74% higher than baseline systems at the same price, highlighting its effectiveness for trace-level detection.

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

Sensys'26, May 2026, Saint-Malo, France

ACM ISBN 978-x-xxxx-xxxx-x/YY/MM…$15.00
https://doi.org/10.1145/nnnnnnn.nnnnnnn

CCS CONCEPTS

• **Human-centered computing** → **Powder adulteration**; **Mobile devices**; • **Computing methodologies** → *Machine learning*.

KEYWORDS

Powder Adulteration, Spectral Imaging and Reconstruction

ACM Reference Format:
Shanwen Chen[1], Haiyan Hu[1], Yinan Zhu[1], Qian Zhang[1]. 2026. FlourSpec: Cost-Effective Spectral Analysis for Trace-Level Detection of Flour Adulteration. In *Proceedings of ACM Conference (Sensys'26).* ACM, Saint-Malo, France, 14 pages. https://doi.org/10.1145/nnnnnnn.nnnnnnn

1 INTRODUCTION

Flour is a fundamental dietary staple, providing over 50% of the daily caloric intake for approximately 2.4 billion people worldwide, particularly in Asia and Africa, where wheat-based products account for 60-75% of dietary protein [12]. This universal reliance, however, renders flour supply chains vulnerable to economically motivated adulteration, where contaminants like corn powder, potato starch, talc, gypsum or chemical oxidants (e.g., benzoyl peroxide) are introduced to increase bulk or mask quality defects [10, 50]. Such practices carry severe economic and public health consequences, such as increased gastric cancer risk at ≥50 ppm benzoyl peroxide exposure [8], and higher prevalence of renal calcification from chronic talc ingestion [22]. The challenge of detecting these high-risk substances is exacerbated by the extremely low detection limits required for flour adulterants, often at parts per million (ppm) levels.

Unfortunately, most current methodologies cannot simultaneously achieve high detection sensitivity and cost effectiveness. Standard laboratory techniques [18, 38, 46] such as HPLC-MS [30] necessitate specialized instrumentation and controlled laboratory environments, rendering them impractical for field deployment. Hyperspectral imaging systems [39] provide nondestructive analysis but require complex optical setups costing over $12,000 with stringent lighting constraints. Conversely, low-cost alternatives have critical limitations; for instance, portable spectrometers can only detect significant compositional changes and often fail to

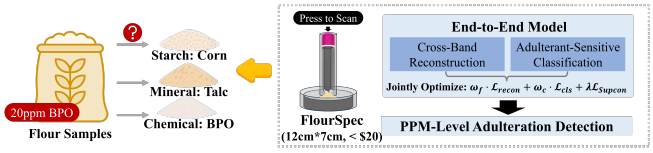

Figure 1: Illustration of FlourSpec 's working scenario.

identify trace adulterants at ppm levels [43]. Electrochemical sensors exhibit reduced sensitivity to nonionic adulterants (such as benzoyl peroxide) due to signal interference [13], while computer vision approaches achieve good accuracy for visible contaminants but struggle with chemical adulterants that exhibit subtle spectral deviations [16]. Consequently, no existing field-deployable system can comprehensively detect multi-class flour adulteration at regulatory-relevant thresholds.

To address this gap, we propose FlourSpec, a low-cost system explicitly designed for ppm-level detection of multiple flour adulterants. Detecting adulterants at ppm concentrations is uniquely challenging: its spectral signatures are weak relative to native flour constituents, partially overlap with flour's broad biochemical bands, and are easily masked by device noise, illumination drift, and sample-to-sample variability in moisture and protein content. Moreover, ppm-level adulterant induces subtle, harmonically related perturbations rather than strong primary peaks, making conventional reconstruction-only approaches prone to false negatives. Prior systems [20, 47] target different matrices and concentration regimes, and thus do not resolve these flour-specific, ultra-trace constraints. Our key contribution is a hybrid attention mechanism that exploits harmonic correlations in the reconstructed spectra to selectively amplify weak adulterant-specific signatures from coarse-grained, low-cost measurements. While spectral reconstruction has been explored to increase granularity [19, 20, 34], coupling it with harmonics-aware attention enables robust, multi-adulterant detection at ppm concentrations using commodity spectrometers.

Firstly, many prior approaches adopt a construction-then-classify cascade structure, which can introduce random errors during the reconstruction phase that interfere with subsequent classification tasks. This is particularly detrimental for detecting low-concentration flour adulteration, as the spectral differences of the adulterants may be obscured by reconstruction errors, making detection nearly impossible. To overcome this challenge, FlourSpec introduces a revolutionary end-to-end structure that enables joint optimization of spectral fidelity and adulteration classification, thereby eliminating the propagation of intermediate reconstruction errors (in § 3.1). To facilitate quick convergence in both tasks, we have designed a learnable dynamic weight adjustment mechanism that prioritizes spectral reconstruction accuracy during the initial training stage, establishing a solid feature foundation, while subsequently shifting focus to classification performance to prevent adulteration-related features from being overwhelmed by reconstruction errors.

Additionally, common flour adulterants, such as talcum and BPO, have characteristic absorption wavelengths in the short-wave near-infrared region (1700-2500 nm), which is beyond the detection range of low-cost spectrometers, typically limited to below 1700 nm. Existing spectral reconstruction algorithms tend to introduce larger errors during cross-band reconstruction, further obscuring effective information from low-concentration adulterants. Although some studies [20] have attempted to use multi-branch feature extractors and knowledge distillation to mitigate reconstruction errors, this multi-branch structure increases the model's parameter count and the risk of overfitting, particularly in non-imaging spectral analysis with limited data. Fortunately, since the fundamental frequency and harmonic frequencies of the same chemical bond exhibit similar absorption characteristics across different bands, adjacent spectral bands demonstrate strong correlations, creating opportunities for effective cross-band reconstruction. Based on this insight, FlourSpec employs a spectral reconstruction structure that combines a multi-layer perceptron (MLP) with a hybrid attention mechanism (in § 3.2); the MLP captures linear reconstruction relationships within the hardware range (400-1700 nm), while the hybrid attention module captures nonlinear harmonic correlations in the range of 1700-2500 nm.

Finally, given the varying regulations surrounding BPO adulteration in different countries, it is essential to identify BPO at various concentration gradients. However, the chemical composition of the same adulterant remains consistent, leading to overlapping absorption peak positions that challenge traditional classification models. To address this, we incorporate a supervised contrastive learning module that aims to reconstruct the feature space to enhance the distinguishability of trace features (in § 3.3). By utilizing cosine similarity between features after L2 regularization as the distance metric for contrast loss, this design enables the model to simultaneously capture differences in absorption peak positions and intensities.

We have developed a fully functional prototype of FlourSpec at a cost of less than $20, as shown in Figure1. The system can identify 10% inexpensive substitutes, including corn, and potato starch, as well as bulking agents like talc and gypsum, and detect BPO at concentrations of at least 20 ppm in both low-gluten and high-gluten flour samples. Results show that FlourSpec achieves 97.44% accuracy in detecting multiple adulteration types and 93.12% accuracy in identifying various BPO concentrations. These results are only 1.27% and 0.63% lower than expensive professional solutions, respectively. Compared to the baseline, our system

is 10.89% higher for adulteration and 18.74% higher for BPO, demonstrating its effectiveness. Moreover, we verified that FlourSpec is robust across various environmental conditions, including temperature, volume, and ambient illumination.

In summary, we make the following contributions:

- We develop FlourSpec, the first ppm-level flour adulteration detection system capable of accurately distinguishing pure flour from various types of adulteration.
- We propose a novel trainable end-to-end spectral reconstruction and classification model that adaptively adjusts training targets to effectively prevent the transmission of reconstruction errors. Additionally, we design a cross-band spectral reconstruction mechanism that leverages harmonic correlations in the spectrum to surpass cross-band reconstruction errors.
- Experimental results demonstrate FlourSpec's high accuracy in detecting various adulteration types and its robustness under different experimental and environmental conditions.

2 BACKGROUND AND CHALLENGES

2.1 Background

Chemical Basis of Flour Adulteration. Flour adulteration remains a pervasive global challenge, with common adulterants categorized into three primary types: **starch-based extenders** (e.g., corn flour, potato starch), **mineral fillers** (e.g., talc, gypsum), and **chemical additives** like BPO [33]. These adultrants will alter the chemical composition of flour, compromising its nutritional integrity and safety profile. Starch-based substitutes, such as corn flour, have distinct biochemical signatures, including higher amylose/amylopectin ratios but lower protein content (6-9% compared to wheat's 10-14%). Inorganic fillers, like chalk and talc, disrupt the flour's physical matrix and lack organic functional groups. Chemical additives like BPO further alter flour composition by oxidizing carotenoids and sulfhydryl groups in gluten proteins [37]. Thus, the chemical compositions of flour samples are an essential clue for identifying adulteration.

Principle of VNIR spectroscopy. Visible-Near Infrared (VNIR) spectroscopy, operating on the principle of photon-matter interactions with specific molecular bonds, can translate various chemical differences in flour into quantifiable spectral patterns across the 400–2500 nm range [50]. Pure wheat flour displays a consistent spectral signature that reflects its balanced chemical composition, while adulterants create measurable deviations from this baseline due to their inherent chemical differences. For instance, corn flour increases absorbance between 400–550 nm due to carotenoid pigments and shifts C–H vibrational bands near between 2100-2200 nm [2]. Chalk suppresses characteristic absorption peaks at approximately 970 nm, 1200 nm, and 1450 nm

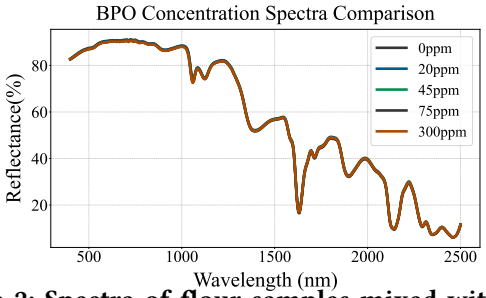

Figure 2: Spectra of flour samples mixed with various concentration of BPO, showing high similarity between various adulteration concentrations.

through scattering-dominated attenuation and exhibits Mg-OH absorption at 2310 nm [17]. BPO decreases absorbance in the visible range (400–500 nm) by bleaching carotenoids and alters N–H band profiles near 1500 nm through protein oxidation [24]. These multivariate spectral variations provide a solid foundation for distinguishing between different adulterants.

Rational of Spectral Reconstruction. Accurate identification of spectral variations in flour adulteration requires fine-grained reflectance measurements $R(\lambda)$, prompting the development of computational reconstruction algorithms that recover high-resolution spectra from sparse affordable spectrometer data. These algorithms utilize inherent spectral redundancy (correlation between neighboring wavelengths) and sparse representation (limited fundamental spectral signatures) to reconstruct dense spectra while maintaining hardware affordability [19, 20, 34, 35]. The process is modeled by relating laboratory-grade measurements $I_{lab} = E(\lambda) \cdot R(\lambda) \cdot S(\lambda)$ to sparse affordable measurements $I_{limited}$, where $E(\lambda)$ and $S(\lambda)$ represent illumination and sensor sensitivity respectively. Reconstruction establishes a transformation matrix W such that $I_{recon} = W \cdot I_{limited}$ approximates I_{lab}.

2.2 Challenges

Despite the effectiveness of computational methods in reconstructing spectra from sparse affordable-device measurements, residual errors compared to laboratory-grade spectra persist. Such inaccuracies critically compromise multi-adulterant detection since identifying trace contaminants like talc or BPO demands high sensitivity to subtle spectral features which are easily obscured by reconstruction artifacts and subsequently degrading classifier performance. These limitations collectively expose three fundamental challenges in achieving robust multi-adulterant detection via reconstructed spectra.

Similarity of Trace Spectral Features. Trace adulterants are characterized by their weak molecular signatures, which can be difficult to detect with cost-constrained sensors. This challenge is particularly evident when examining the

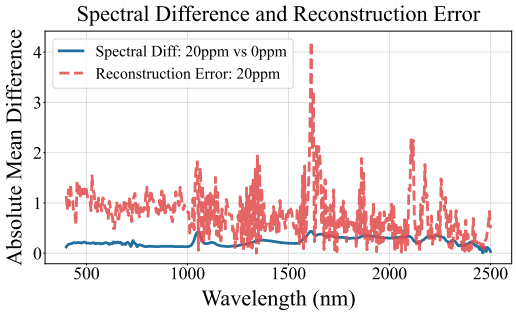

Figure 3: The spectra difference between a pure flour sample and a 20 ppm BPO adulterated sample, and the distribution of reconstruction errors for the 20 ppm BPO adulterated sample using the SOTA reconstruction method.

spectral data for BPO, as illustrated in Figure 2. The figure highlights the spectral variations across concentrations that are relevant to regulatory standards, reflecting international safety thresholds: China (≤60 ppm, banned since 2011), USA (≤50 ppm), Codex (≤75 ppm), Japan (≤300 ppm) [41, 42]. Despite these varying thresholds, the spectral differences observed at these concentrations are minimal, indicating a high degree of similarity among the spectra. This pronounced similarity complicates the task of distinguishing chemical fingerprints, making it increasingly difficult to identify trace adulterants based solely on spectral data.

Obscuration from Reconstruction Errors. The reconstruction process further complicates the detection of trace adulterants by introducing errors that are randomly distributed across all spectral bands of the reconstructed data. Given the subtle spectral characteristics of trace elements, these reconstruction errors can obscure critical spectral information related to adulterants, ultimately leading to degraded prediction performance. Figure 3 illustrates the spectral differences between a pure flour sample and a 20 ppm BPO-adulterated sample, alongside the distribution of reconstruction errors for the adulterated sample, which was reconstructed using a state-of-the-art algorithm [19]. The figure clearly demonstrates that the reconstruction process contaminates the absorption peak information of the raw spectra, as the magnitude of the reconstruction error exceeds the differences between the samples. This highlights the importance of extracting useful features during the reconstruction process, as it is essential for accurate detection of flour adulteration.

Limitations of Spectral Coverage. As discussed in § 2.1, the characteristic bands of common flour adulterants span a broad range, typically from 400 to 2500 nm [50]. While the visible and near-infrared bands (400-1700 nm) contain some discernible information, the data presented in Figure 4 indicate that relying solely on spectral data from this range

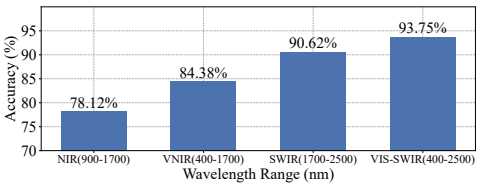

Figure 4: Classification performance using different range of light channles using SOTA CNN model.

is inadequate for detecting low-concentration flour adulteration. This limitation is particularly pronounced for chemical adulterants like BPO, whose primary spectral absorption characteristics occur above 1700 nm. Unfortunately, the high cost of optical sensors that operate beyond 1700 nm restricts their availability for low-cost near-infrared spectrometers, which typically only provide LEDs within the lower range. In the MeatSpec project [20], a multi-branch feature extractor was developed to address errors in cross-band spectral reconstruction. However, this multi-branch architecture increases the model's parameter count, thereby raising the risk of overfitting, especially when working with one-dimensional spectra that have limited data availability. Consequently, it is crucial to explore alternative low-consumption approaches for cross-band reconstruction to enhance detection capabilities without compromising model performance.

3 SYSTEM DESIGN

To solve the challenges mentioned above, we propose FlourSpec which contains three novel designs to obtain flour adulteration detection from affordable LED-array spectrometer, as shown in Figure 5.

- **Firstly**, we redesign the conventional reconstruction-then-analysis pipeline into an **end-to-end trainable framework** (§ 3.1), enabling joint optimization of spectral fidelity and adulteration classification while eliminating intermediate reconstruction error propagation.
- **Secondly**, We design a **hybrid attention module** to extract local spectral features related to flour adulteration in the reconstructed spectrum, and simultaneously utilize the neighboring correlation between spectral bands to achieve reconstruction of the extended band, *i.e.*, 400-2500nm (§ 3.2).
- **Thirdly**, we employ a **supervised contrastive learning module**, with the core objective of restructuring the feature space such can enhance the distinguishability of trace spectral features (§ 3.3).

3.1 End-to-End Architecture

As we discussed in § 2.2, traditional cascade architectures [20, 34] suffer from error amplification, where reconstruction errors directly impair contaminant detection, as critical spectral features are distorted before reaching the classifier. Thus, to mitigate the intermediate reconstruction error, FlourSpec

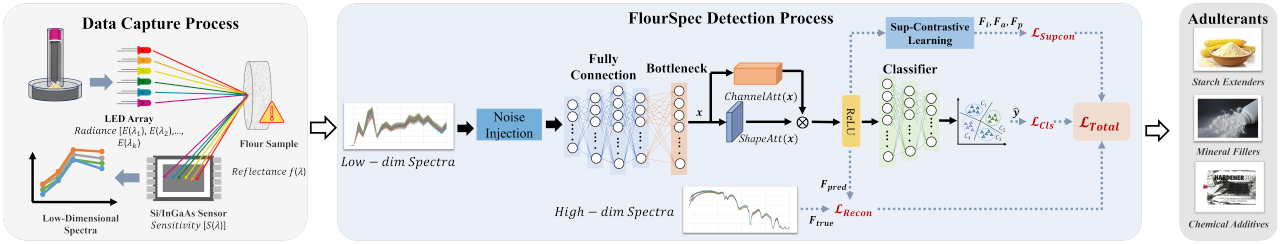

Figure 5: The architecture of FlourSpec, which is an end-to-end framework that jointly optimize the spectral fidelity and adulteration classification.

proposes a new end-to-end structure with integrated reconstruction and classification process. In this way, we can guide the model to reconstruct the most flour-related spectral feature instead of obscurating the spectral feature by random reconstruction errors.

Dynamic Loss Balancing Mechanism. However, we find that directly combining the reconstruction module and the classification module together would make it difficult for the model to converge, making it difficult to effectively play the role of the two models. To this end, FlourSpec utilizes a dynamic loss balancing mechanism to balance the reconstruction and classification objectives simultaneously. Specifically, the joint objective function is defined as following:

$$\mathcal{L}_{\text{joint}} = \underbrace{\text{softmax}(\boldsymbol{\omega}_f) \cdot \mathcal{L}_{\text{Recon}}}_{\text{spectral fidelity}} + \underbrace{\text{softmax}(\boldsymbol{\omega}_c) \cdot \mathcal{L}_{\text{Cls}}}_{\text{contaminant detection}}, \quad (1)$$

where $\mathcal{L}_{\text{Recon}}$ denotes the mean squared error between reconstructed and true spectra, $\mathcal{L}_{\text{Cls}}$ is the classification loss. To dynamically balance the two objectives, we introduce a trainable parameter vector $\boldsymbol{\omega} = [\omega_f, \omega_c]$, initialized as $[0.7, 0.3]$. During training, these weights dynamically evolve. In the early epochs, ω_f increases to emphasize spectral reconstruction, while in later epochs, ω_c increases to enhance classification performance. This adaptive weighting ensures that spectral reconstruction accuracy is prioritized in the initial training phase to build a solid feature foundation, while subsequent phases shift the focus to classification performance for adulteration detection to avoid adulteration-related features being overwhelmed by reconstruction errors.

Noise-Robust Training Strategy. Additionally, to further suppress the random errors introduced by the model during the reconstruction process, we employ a comprehensive noise-robust training strategy at both data and feature levels. At the data level, Gaussian noise is injected into training samples with a standard deviation of $0.1\sigma_{\text{data}}$ and a probability of 10%, i.e., $x' = x + \mathcal{N}(0, 0.1^2\sigma^2_{\text{data}})$, which simulates random fluctuations and improves the model's ability to combat random noises. At the feature level, We insert batch normalization layers after each linear transformation to stabilize feature distributions, thereby further enhancing the model's resistance to input perturbations [31].

Additionally, we incorporate an implicit confidence weighting mechanism to penalize excessive local reconstruction errors. Specifically, features corresponding to regions with high reconstruction error are automatically down-weighted through the gradient flow of the reconstruction loss:

$$\nabla_\theta, \mathcal{L}\text{recon} \propto F_{\text{pred}} - F_{\text{true}}, \quad (2)$$

which ensures that samples with larger errors receive stronger gradient updates. This mechanism guides the model to focus on more challenging areas without the need for explicit confidence computation, further boosting robustness under real-world noise conditions.

3.2 Cross-Band Spectral Reconstruction

Flour adulteration requires a broad absorption wavelength range, spanning visible light (400-700 nm), near-infrared (700-1700 nm), and short-wave infrared (1700-2500 nm). However, low-cost near-infrared spectroscopy equipment only supports the visible-near infrared range (400-1700 nm). Therefore, the spectral reconstruction module of the system needs to have the ability to reconstruct across bands. Unlike the multi-branch feature extractors used in [20] to address cross-band reconstruction errors during classification, FlourSpec aims to eliminate these errors in the reconstruction stage to prevent their impact on classification. As shown in Figure 5, the FlourSpec spectral reconstruction module includes a multilayer perceptron (MLP) for linear reconstruction and a hybrid attention module for nonlinear reconstruction.

Multilayer Perceptron Module. For the bands covered by the hardware device (*i.e.*, 400-1700 nm), the goal is to learn the linear mapping matrix of the hardware device parameters. Therefore, we adopt the design logic from [19] and implement a multi-layer MLP structure to achieve the linear mapping from the low-dimensional spectrum (44) to the high-dimensional spectrum (431). However, unlike the approach in [19], which combines linear regression with a denoising autoencoder, FlourSpec employs a direct 4-layer MLP structure (44[input] $\rightarrow$ 256 $\rightarrow$ 512 $\rightarrow$ 256 $\rightarrow$ 431) as a replacement. This choice is based on two reasons. First, both

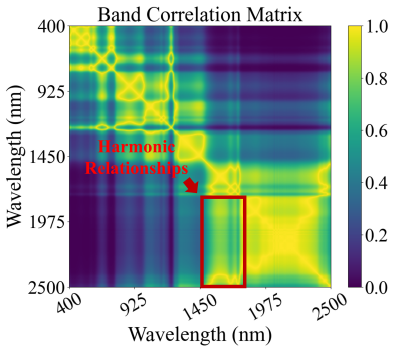

Figure 6: The correlation matrix between the spectral bands of 400-2500 nm of a flour sample.

the input and output dimensions of FlourSpec are four times larger than those in [19], making it challenging to effectively capture the complete correlation using simple linear regression. Second, FlourSpec utilizes an end-to-end optimization strategy, where the final joint loss is used to optimize the entire reconstruction-regression process for optimal detection performance. The use of linear regression would hinder the effective optimization of the linear reconstruction model by subsequent processes.

Hybrid Attention Module. For the bands outside the coverage range of the hardware device (*i.e.*, 1700-2500 nm), traditional linear-based reconstruction solutions tend to introduce larger reconstruction errors in these extended bands, as validated in [20]. Consequently, achieving complete cross-band reconstruction using only a single MLP structure is challenging. Fortunately, we have observed that adjacent spectral bands exhibit strong correlations, which may provide opportunities for effective cross-band reconstruction. Figure 6 illustrates the correlation matrix between the spectral bands of 400-2500 nm for a flour sample, revealing a strong correlation between the spectra in the 1450-1700 nm range and those in the 1700-2500 nm range. This correlation arises from the harmonic relationships of molecular vibrations, where overtones of the same chemical bonds manifest in adjacent spectral regions. For example, the first overtone of C-H bond stretching occurs near 1700 nm, while the third overtone appears around 2200 nm. These harmonic relationships establish a physical basis for cross-band feature transfer, inspiring us to address the errors of cross-band reconstruction by leveraging the correlations between bands.

To achieve this goal, we design a hybrid attention module that is integrated after the MLP module. This module extracts local spectral features related to flour adulteration in the reconstructed spectrum while simultaneously utilizing the neighboring correlations between spectral bands to facilitate the reconstruction of the extended band. This mechanism comprises two complementary branches: channel attention and shape attention, which are followed by an attention fusion module.

Channel Attention. The channel attention branch is responsible for learning the importance of each spectral band, especially those associated with key harmonic relationships. This branch utilizes a squeeze-and-excitation structure: the input feature $x \in \mathbb{R}^{B\times D}$ (where B is the batch size and $D = 421$ is the feature dimension) is first compressed using a learnable weight matrix $\mathbf{W}_1 \in \mathbb{R}^{D\times D/8}$, activated by ReLU, then expanded back by $\mathbf{W}_2 \in \mathbb{R}^{D/8\times D}$ and mapped through a sigmoid function to produce channel attention weights:

$$\text{ChannelAtt}(x) = \sigma\left(\mathbf{W}_2\delta(\mathbf{W}_1 x)\right), \tag{3}$$

where $\delta(\cdot)$ denotes the ReLU activation and $\sigma(\cdot)$ denotes the sigmoid function. This process allows the model to assign greater significance to spectral bands that are critical for accurate reconstruction, particularly those involved in cross-band harmonic relationships.

Shape Attention. The shape attention branch aims to capture the continuity between adjacent spectral bands, which is essential for reconstructing smooth spectral curves. This branch applies a one-dimensional convolution with kernel size $k = 5$ and padding $p = 2$ to the input, followed by a sigmoid activation to produce shape attention weights:

$$\text{ShapeAtt}(x) = \sigma\left(\text{Conv1D}_{k=5,p=2}(x)\right) \tag{4}$$

This operation enables the model to capture local modes in the spectrum, such as absorption peaks and their surrounding areas, ensuring that the reconstructed spectra maintain local continuity. Since harmonics are essentially integer multiples of bond vibrations, the spectrum reveals that the fundamental frequency and its second harmonic exhibit similar yet broadened peak shapes. By establishing this local receptive field, the model can explicitly model these harmonic relationships, allowing it to transfer characteristic peak patterns from covered fundamental vibrations (*e.g.*, C-H bonds in the 1600-1700 nm range) to reconstruct the corresponding overtones (*i.e.*, 2000-2200 nm) in uncovered regions .

Attention Fusion. To exploit the strengths of both channel and shape attention, we introduce a fusion mechanism. Specifically, the channel attention weights (reflecting the importance of each band) and the shape attention weights (capturing local continuity) are combined by element-wise multiplication, employing broadcasting to align their shapes. The fused attention is then applied to the feature representations as:

$$F_{\text{att}} = x \circ \left(\text{ChannelAtt}(x) \otimes \text{ShapeAtt}(x)^T\right) \tag{5}$$

where $\circ$ denotes element-wise multiplication and $\otimes$ denotes the broadcasting operation. In implementation, channel attention weights are reshaped to $[B, D, 1]$, shape attention

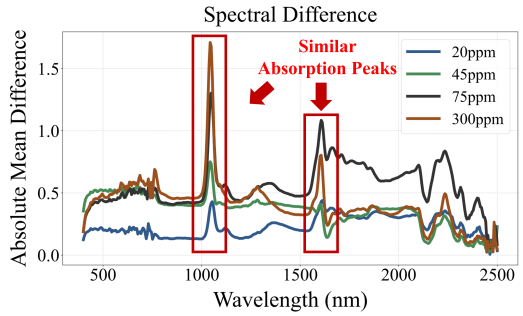

Figure 7: Spectra differences between a pure flour sample and various concentration of BPO adulterated samples.

weights to $[B, 1, D]$, and after broadcasting, the resulting attention matrix $[B, D, D]$ is applied to the feature tensor.

With this dual-attention mechanism, FlourSpec achieves high-fidelity reconstruction of the full spectrum from narrowband LED signals while maintaining real-time inference speed. This design effectively bridges the information density gap and captures essential harmonic relationships, providing a reliable spectral foundation for ppm-level contaminant detection in practical deployments.

3.3 Adulterant-Sensitive Classification

Based on the aforementioned cross-band reconstruction architecture, we then use a two-layer MLP (421[input]->128->N) for multi-classification of flour adulteration, where N refers to the number of output types. It is worth noting that FlourSpec uses a fully connected layer instead of the convolutional neural network-based structure commonly used in traditional spectral analysis [17, 19, 20]. This is mainly because MLP has stronger global modeling capabilities and can directly capture the harmonic correlation features across bands in the reconstructed spectrum, while the local receptive field of CNN will dilute such long-range dependencies. Especially for the flour adulteration task, each adulterant may involve a combination of features from multiple bands. For example, corn flour requires the absorption features brought by carotenoids in the visible light band and the absorption features caused by the C-H bonds of starch in the near-infrared region. Therefore, MLP, a modeling that has the ability to perceive the global spectrum, is very necessary.

Accurate flour adulteration detection faces significant challenges due to spectral similarities between contaminants, particularly at low concentrations. As shown in Figure 2, pure flour and different BPO concentrations exhibit minimal spectral variance, complicating fine-grained discrimination. To overcome this limitation, FlourSpec employs a supervised contrastive learning module, with the core objective of reformulating the feature space such that samples of the same adulterant concentration or type are pulled together while those of different concentrations or types are pushed apart. Different from the supervised contrastive learning used in [20] that utilize Euclidean distances to measure the similarity between samples, we utilize the cosine similarity on ℓ_2-normalized features to measure their similarity, which shows more effective for flour adulteration task. Specifically, we find that the the same substance at different concentrations may have similar absorption peak shapes, but the absorption intensity changes linearly with concentration. For example, adulterated samples of 20 ppm BPO and 300 ppm BPO may have the same absorption peak shape due to the same chemical bond type, but with very different amplitudes, as shown in Figure 7. This may cause the model classification to be dominated by amplitude differences while ignoring the differences in absorption peak shapes, which is not conducive to the joint detection of multiple types of adulterants. To this end, all feature vectors are ℓ_2-normalized to remove amplitude effects and focus solely on angular information, leading to $\tilde{\mathbf{f}} = \frac{\mathbf{f}}{\sqrt{\sum_{k=1}^{421} f_k^2}}$. Then we input the normalized features into supervised contrastive loss calculation, which is defined as:

$$\mathcal{L}_{\text{Supcon}} = -\frac{1}{B}\sum_{i=1}^{B}\frac{1}{|\mathcal{P}(i)|}\sum_{p\in\mathcal{P}(i)}\log\frac{\exp(\cos(\tilde{\mathbf{f}}_i,\tilde{\mathbf{f}}_p)/\tau)}{\sum a\neq i\exp(\cos(\tilde{\mathbf{f}}_i,\tilde{\mathbf{f}}_a)/\tau)} \tag{6}$$

where B is the batch size, $\mathcal{P}_i$ denotes the set of positives for anchor i (i.e., samples with the same concentration or class), $\cos(\cdot)$ is the cosine similarity, and τ is the temperature parameter. Since we find that the cosine similarity between low-concentration BPO and pure wheat is smaller, we set a small temperature value (empirically set to 0.07) to prioritize low-concentration detection.

By enforcing intra-class compactness and inter-class separation in the learned spectral feature space, the model markedly improves its ability to resolve subtle spectral differences introduced by low-level adulteration, including BPO, various starches, and bulking agents. This enables robust multiclass discrimination across all tested concentrations and adulterant types, directly addressing the challenge of sensitive and specific flour adulterant detection. Overall, the loss function is formulated as $\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{joint}} + \lambda \cdot \mathcal{L}_{\text{Supcon}}$, where λ is set to control the weight of contrastive loss (empirically set to 0.3). Consequently, by jointly optimize this combined loss function, FlourSpec can achieves accurate flour adulteration detection from low-cost spectrometer input.

4 IMPLEMENTATION

FlourSpec incorporates a commercial Vispek Magic Pencil spectrometer (26 × 99.5 mm, 65 g) housed within a custom 3D-printed enclosure made from light-absorbing black PLA resin, as shown in Figure 8. The per-device hardware cost (Magic Pen used in our work; internal pricing) is under $15.

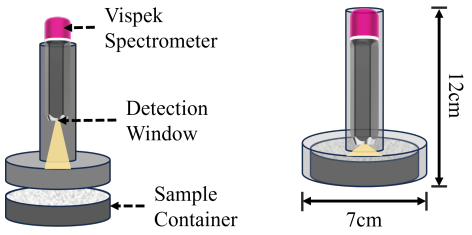

Figure 8: Prototype of FlourSpec,which consists of a commercial Vispek spectrometer and a enclosure. Users can finish adulteration detection by simply press the button at the bottom.

Our 3D-printed enclosure adds approximately $5. Therefore, the total device cost is about $20. The system is designed for single-button operation, allowing users to initiate spectral acquisition and capture complete data within a rapid 500 ms cycle.

Enclosure Design and Functionality. The optically sealed enclosure, designed in Tinkercad and fabricated using fused deposition modeling at a 0.1 mm layer resolution, features a sample chamber with a diameter of 70 mm and a height of 15 mm, along with an integrated spectrometer mount. When sealed, this configuration keeps ambient light levels below 1 lux and precisely aligns the spectrometer's fixed-focus probe (9 mm spot size, 90° incidence angle) with the sample surface. This design effectively eliminates environmental interference and ensures measurement repeatability, even for untrained operators.

Spectrometer Selection and Parameters. The choice of the Vispek spectrometer is based on its exceptional spectral coverage, ranging from 200 to 1700 nm[1]. This range significantly exceeds that of comparable devices, such as the AS7421 (380-940 nm) [29] and the SCiO Mini (740-1070 nm) [32]. The extensive spectral range includes critical diagnostic bands for flour adulterants, such as the carbonyl peak of benzoyl peroxide found near 1700 nm. Furthermore, the Vispek spectrometer boasts an optical resolution of 33 nm, ensuring the high-quality raw data necessary for effective spectral reconstruction. In terms of performance, Vispek surpasses its competitors in sensitivity, achieving 0.96 A/W at 940 nm on the silicon sensor and 0.95 A/W at 1550 nm on the InGaAs sensor. Additionally, it demonstrates excellent wavelength repeatability, with less than 1% variation in the visible and infrared ranges. These features make the Vispek spectrometer particularly well-suited for trace chemical detection while maintaining cost-effectiveness.

[1]The Vispek spectrometer incorporates 22 different wavelength LEDs (255nm, 275nm, 370nm, 400nm, 455nm, 520nm, 560nm, 610nm, 630nm, 667nm, 700nm, 804nm, 850nm, 890nm, 933nm, 973nm, 1000nm, 1217nm, 1300nm, 1470nm, 1545nm, 1635nm), each providing two sets of readings by applying high and low voltage levels, resulting in a comprehensive 44-dimensional output.

Table 1: Adulterant types and their concentrations in our flour adulteration dataset.

Types	Adulterants	Concentration
Starch	Corn flour	10%, 50%
	Potato starch	10%, 50%
Mineral	Talc	10%, 50%
	Gypsum	10%, 50%
Chemical	Benzoyl peroxide (BPO)	20 ppm, 50 ppm, 75 ppm, 300 ppm

5 PERFORMANCE EVALUATION

5.1 Study Setup

5.1.1 Data Collection. We collected commercial flour samples from the QUEEN brand, including both high-gluten and low-gluten varieties (see Figure 9(a)). For each flour type, we prepared samples with and without adulterants. Table 3 summarizes the adulterants in our dataset: corn powder, potato powder, talc, gypsum, and BPO. For the first four adulterants, samples were prepared at two levels (10% and 50%), with six samples for each level and flour type. For BPO, samples were prepared at four concentrations (20, 50, 75, and 300 ppm), with six samples per concentration and flour type. Additionally, the extended BPO dataset includes 16 samples for each concentration and 16 pure flour samples for both flour types. In total, we prepared **316** flour samples: the first dataset contains 12 pure flour samples, 24 samples each adulterated with corn powder, potato powder, talc, and gypsum, and 48 BPO samples; the extended BPO dataset adds 32 pure flour samples and 128 BPO samples (32 per concentration).

All samples were scanned once in custom 3D-printed black plastic holders using our prototype device, yielding a 44-dimensional low-resolution spectrum per sample. Ground truth high-resolution spectra (431 dimensions, 400–2500 nm, 5 nm resolution) were collected with a PerkinElmer LAMBDA 1050+ UV/Vis/NIR spectrometer using an integrating sphere. All experiments were conducted at 25–28°C under standard office lighting (100–250 lx).

5.1.2 Metrics. We reformulate the flour adulteration detection task as two multi-class classification problems, which involve not only distinguishing pure flour from adulterated samples, but also identifying the specific adulterant type or, in the case of BPO, different adulterant concentrations. For the first task, we consider five common adulterants—corn powder, potato powder, talc, gypsum, and BPO—as well as pure flour, resulting in six classes: pure flour (0), corn powder (1), potato powder (2), talc (3), gypsum (4), and BPO (5). For the second task, we focus on classifying samples with different concentrations of BPO (20 ppm, 50 ppm, 75 ppm, and 300 ppm) and pure flour, forming a five-class task: pure flour (0), 20 ppm BPO (1), 50 ppm BPO (2), 75 ppm BPO (3),

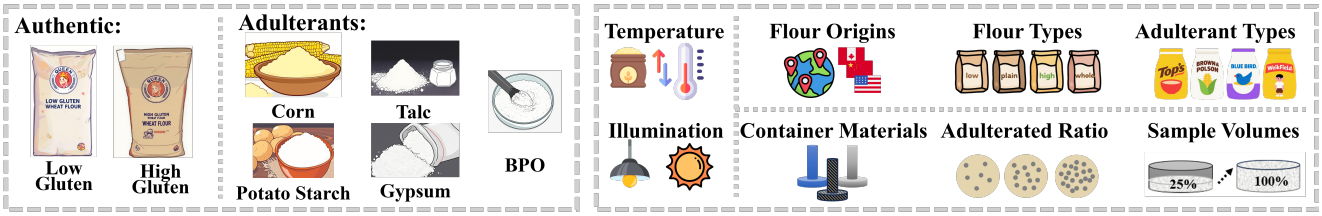

(a) Overall Samples (b) Various Experimental Conditions

Figure 9: The setup of the evaluation. (a) We use two types of authentic wheat flour and five common adulterants to evaluate the FlourSpec. (b) We validate the robustness of FlourSpec across eight different experimental conditions.

and 300 ppm BPO (4). **We will make our datasets open source after the review process.** It should be noted that, for both classification tasks, samples from both high-gluten and low-gluten flour are pooled together for modeling and analysis. To evaluate the performance of the multi-class flour adulteration classification tasks, we adopt **macro-averaging accuracy** and **F1 score** as our evaluation metrics.

5.1.3 Baselines. Since there is no low-cost automatic flour adulteration detection solution, we consider to apply the SOTA spectral reconstruction algorithm and CNN-based spectral analysis model[2] as baselines instead, similar to the works of other foods [6, 19]. Besides, we also take the performance of affordable 44D spectrometer without spectral reconstruction as a baseline. The details are as follows:

- Raw 44D spectra + CNN. We directly use the 44-dimensional low-resolution spectral data for classification, and employ a CNN classifier for this task.
- Reconstructed 431D spectra + CNN [19]. This baseline mimics the traditional cascade solution. We use the state-of-the-art spectral reconstruction in [19] to process the 44D spectra and conduct classification using the CNN classifier for classification.
- GT 431D spectra + CNN. We use the ground truth high-dimensional spectra as input to the CNN classifier for classification as the third baseline. This baseline is expected to have better performance than FlourSpec, due to the high quality and cost of the hardware.

5.2 Overall Performance

We first evaluate FlourSpec's performance in detecting different types of adulteration by conducting stratified 5-fold cross-validation.

5.2.1 Baselines Comparison. Table 2 compares our FlourSpec system with several baselines under various conditions and highlights three key findings. First, directly using the 44-dimensional low-resolution spectra with a CNN (44D+CNN) results in limited accuracy for both BPO concentration (65.62%) and adulterant type detection (78.12%), due to insufficient spectral detail. Second, the state-of-the-art BabyNutri [19] model improves results (74.38% for BPO and 86.55% for adulterant type), but still falls short in capturing subtle spectral variations. Most notably, FlourSpec achieves much higher accuracy—93.12% for BPO concentration and 97.44% for adulterant type—approaching the upper bound of the GT+CNN baseline (93.75% and 98.71%). These results demonstrate that FlourSpec can reliably reconstruct high-fidelity spectra from low-resolution data, enabling robust and accurate adulteration detection.

5.2.2 FlourSpec Performance and Reconstruction Fidelity. Figure 10 shows the confusion matrices of FlourSpec for

[2]We have tested several common spectral analysis algorithms, including PLS, SVM, and MLP, but find that CNN performs the best in this task. Its structure aligns with an open-source project at https://github.com/FuSiry/OpenSA.

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

Scenario	Models	Accuracy	F1-Score
BPO concentration	44D+CNN	65.62%	64.05%
	Recon+CNN	74.38%	74.02%
	FlourSpec	93.12%	92.95%
	GT+CNN	93.75%	93.72%
Adulterants Type	44D+CNN	78.12%	77.44%
	Recon+CNN	86.55%	80.30%
	FlourSpec	97.44%	96.05%
	GT+CNN	98.71%	97.55%

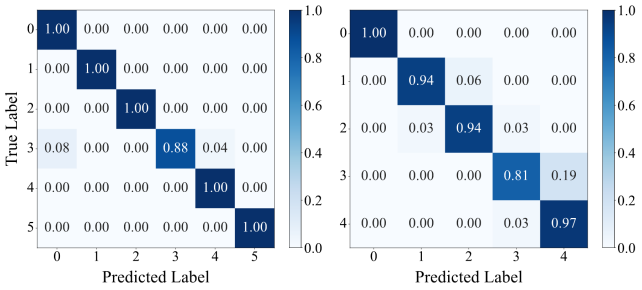

(a) Adulterants Type (b) BPO concentration

Figure 10: Confusion matrix of FlourSpec. (a) Confusion matrix on multiple Flour adulteration types classification. (b) Confusion matrix on different BPO concentration detection.

Table 3: FlourSpec 's performance compared with ablation models and SOTA reconstruction models.

Model	Accuracy	F1-Score
SOTA	74.38%	74.02%
FlourSpec (w/o DA)	86.31%	86.25%
FlourSpec (w/o ASC)	86.88%	86.79%
FlourSpec (ours)	93.12%	92.95%

both flour adulteration type detection and BPO concentration classification. As seen in Figure 10 (a), FlourSpec consistently achieves high accuracy across all adulteration types, with most categories being precisely identified. Similarly, Figure 10 (b) demonstrates that FlourSpec accurately distinguishes different BPO concentrations, with only minor misclassifications for closely related classes.These results highlight FlourSpec 's strong capability in both detecting diverse adulterants and identifying varying BPO levels in flour samples. We further evaluated the fidelity of FlourSpec's spectral reconstruction by comparing the reconstructed spectra against ground-truth laboratory measurements. Our approach achieves high reconstruction accuracy, with an average RMSE of 0.038 across the full spectral range (400-2500nm). This outperforms the SOTA method [19] (RMSE ≈0.075), ensuring the subtle spectral features indicative of trace-level adulteration are preserved for subsequent classification.

5.3 Ablation Study

We then investigate the modules of FlourSpec by conducting ablation study to demonstrate the effectiveness of the system design. The evaluation is performed on the classification task of distinguishing flour samples with different concentrations of BPO, in order to assess their effectiveness in detecting trace elements.

5.3.1 Effectiveness of Dual-Attention. As shown in Table 3, removing the Dual-Attention module from FlourSpec leads to a significant drop in both accuracy and F1 score (from 93.75%/93.72% to 86.31%/86.25%). This demonstrates that the Dual-Attention module is crucial for restoring important adulteration-related features in spectral regions beyond the device's native range. By leveraging both linear MLP mapping and hybrid attention, our method effectively captures harmonic relationships among adjacent bands, enabling accurate reconstruction of extended spectrum and directly boosting downstream adulterant detection performance.

5.3.2 Effectiveness of Adulterant-Sensitive Classification. Table 3 further shows that removing the adulterant-sensitive classification (ASC) module results in a clear decrease in performance (accuracy drops from 93.75% to 86.88%). This highlights the importance of our classification design, which utilizes a global MLP structure to fully exploit cross-band features and subtle adulterant signatures. The ASC module ensures that discriminative spectral cues are well captured across all adulteration types, leading to significant improvements in both classification accuracy and robustness compared to baseline and SOTA methods.

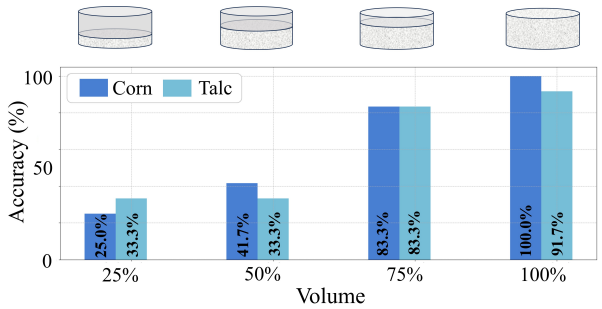

Figure 11: Accuracy with varying volume and height.

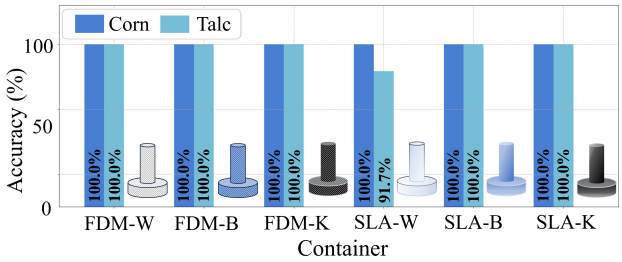

Figure 12: Accuracy with varying container color and material.

5.4 Impacts on Experimental Conditions

We evaluated the robustness of FlourSpec under a variety of experimental and environmental conditions, as illustrated in Figure 9 (b). Focusing on two typical adulterants—corn (starch-based) and talc (mineral-based)—we prepared 12 samples each with 10% adulteration for every condition. All models were trained using data collected under standard conditions.

5.4.1 Varying flour volume and Height. To test the effect of sample filling, we compared accuracy across four filling ratios (25%, 50%, 75%, and 100%; see Figure 11). Classification accuracy improved as the filling ratio increased, with the best results at 100%. Lower ratios reduce illumination and signal collection efficiency, leading to lower accuracy. Thus, adequate sample filling is important for optimal system performance.

5.4.2 Varying container color and material. We tested different 3D-printed sample holders, including FDM and SLA methods and three colors (Figure 12). Classification accuracy remained at 100% across all holder types and colors, demonstrating that the system is robust to changes in holder material and color.

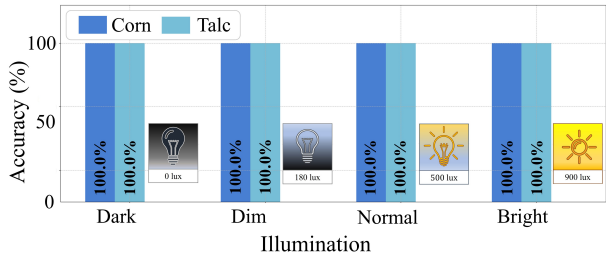

Figure 13: Accuracy with varying ambient illumination.

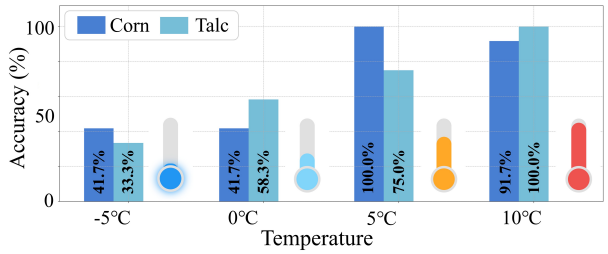

Figure 14: Accuracy with varying ambient temperature.

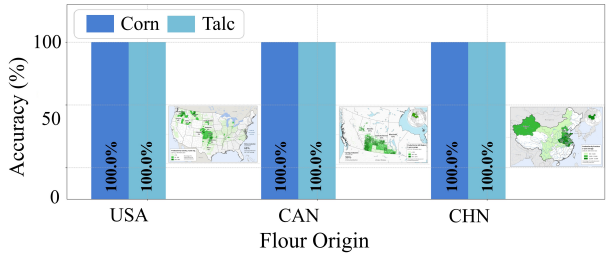

Figure 15: Accuracy with varying authentic flour origin.

5.4.3 Varying Ambient Illumination. The system was evaluated under illumination levels of 0, 180, 500, and 900 lx. As shown in Figure 13, classification accuracy was consistently 100% under all lighting conditions. This is due to the fully enclosed design, which isolates the measurements from ambient light.

5.4.4 Varying Ambient Temperature. Samples frozen at −5°C and then measured at different temperatures showed a sharp drop in accuracy at subzero temperatures (Figure 14), mainly because ice formation alters spectral features. As temperature increases and water returns to liquid, accuracy recovers, indicating temperature stability above freezing.

5.4.5 Varying Authentic Flour Origin. We tested flours sourced from the US, Canada, and China, each adulterated with 10% corn or talc. Classification accuracy was 100% across all origins (Figure 15), indicating that the system is robust to flour geographic origin.

5.4.6 Varying Authentic Flour Type. We evaluated high-gluten, plain, low-gluten, and whole wheat flours adulterated with 10% corn or talc (Figure 16). Accuracy was 100% for most types, except for corn adulteration in plain flour (91.7%) and talc in whole wheat (83.3%). This moderate decrease in performance for talc detection in whole wheat flour may be attributed to the complex matrix interactions between talc particles and the bran components in whole wheat, which could alter the spectral characteristics.

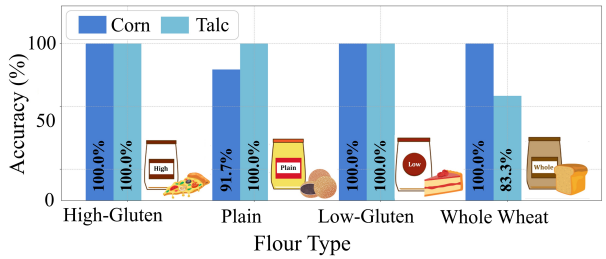

Figure 16: Accuracy with varying authentic flour type.

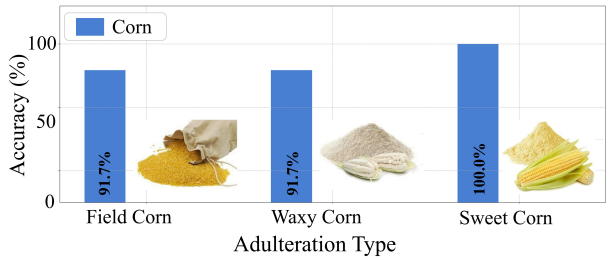

Figure 17: Accuracy with varying adulteration type.

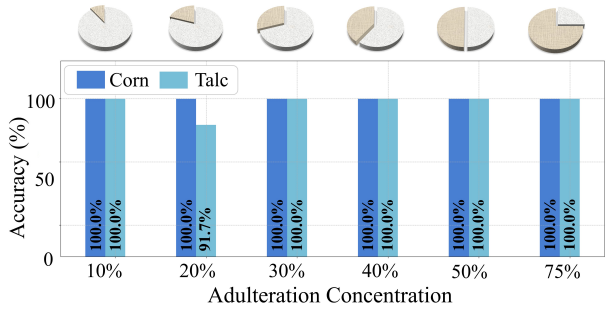

Figure 18: Accuracy with varying adulteration concentration.

5.4.7 Varying Adulteration Type. Three corn flour types—feed, waxy, and standard edible—were tested (Figure 17). Accuracy was 100% for edible corn flour and 91.7% for feed and waxy types, with rare misclassification as potato flour, likely due to similar starch content. Talc was not tested for type variation due to its simple composition.

5.4.8 Varying Adulteration Concentration. We further varied corn and talc adulteration concentrations from 10% to 75% (Figure 18). Corn adulteration was detected at 100% accuracy at all levels; talc accuracy was 100% except for a small drop to 91.7% at 20%. This demonstrates high sensitivity and robustness to different adulteration levels.

6 RELATED WORK

In this section, we briefly review related works on flour adulteration detection and spectral reconstruction.

Flour Adulteration Detection. Traditional methods such as chromatography [1, 4, 9, 15, 30] and molecular biology techniques [18, 38, 46] offer quantitative analysis but require costly equipment and expertise. Spectroscopic approaches including vibrational [7, 51], electronic [23, 25], imaging [39, 45], and NMR spectroscopy [5]—enable non-destructive, rapid detection, but remain too expensive for routine consumer use. Vision-based methods [27, 28] are affordable and portable but limited to visually distinguishable adulteration. In contrast, FlourSpec achieves ppm-level detection of diverse adulterants using sub-$50 commercial spectrometers, bridging the gap between sensitivity and affordability.

Spectral Reconstruction Algorithms. Spectral reconstruction enables high-precision analysis from sparse measurements, supporting low-cost deployment. Classical NIR reconstruction methods—such as linear regression [11], partial least squares [44], and SVM [48]—assume linear relationships, limiting their robustness to environmental variation. Handcrafted prior-based hyperspectral methods [3, 14, 26] face similar limitations. Recent deep learning methods [36, 49] improve mapping but often overlook practical deployment constraints. FlourSpec adopts a lightweight, application-specific architecture, enabling robust flour adulteration detection in daily scenarios.

Low-cost Adulteration Detection Systems. Recent works explore affordable product authentication. LiquidHash [40] and Vi-Liquid [21] analyze physical properties to detect liquid fraud, but cannot handle powder adulteration. MobiSpectral [34] uses smartphone infrared sensors but lacks necessary wavelength coverage for powders. MeatSpec [20] detects meat adulteration but is not suited for lightweight, consumer use. PowDew [47] leverages droplet motion analysis on powders, but cannot reliably detect trace additives. By contrast, FlourSpec advances the state of low-cost adulteration detection systems by specifically addressing the challenges of flour authentication, including the need for broad spectral coverage and sensitivity to trace contaminants.

7 DISCUSSION

In this section, we discuss the limitations and potential extensions of FlourSpec.

Generalizability and Dataset Scope. While FlourSpec performs well on common adulterants, its evaluation is limited to one flour brand (QUEEN) and five adulterant types using lab samples, lacking real-world variability. Future work will expand datasets across brands, regions, and field samples. Methodologically, we will explore transfer learning, domain adaptation, and self-supervised pretraining for better generalization. Uncertainty estimation and out-of-distribution detection will enhance confidence quantification and flag low-confidence samples.

Sampling Depth and Optical Coverage. Trace-level detection remains challenging due to limited optical penetration and field of view, causing false negatives with non-uniform contamination. Solutions include multi-point, multi-depth sampling and result compositing. Hardware improvements like wider field view and optimized illumination geometry will stabilize measurements. Standardized procedures and calibration will improve operational repeatability.

Integration with Other Applications. FlourSpec's principles can extend to other foods (e.g., grains, dairy, oils) via matrix-specific calibration and spectral libraries. Cross-product transfer learning and multi-task modeling will reduce data demands for new categories, while active learning guides efficient data collection. Integration into supply chains requires expanded datasets and tailored protocols for unified food safety surveillance.

8 CONCLUDING REMARKS

This paper presents FlourSpec, an innovative system designed for the detection of flour adulteration at the parts per million (ppm) level. FlourSpec effectively combines low-cost hardware with advanced spectral reconstruction algorithms to ensure accurate identification of common adulterants such as talc and benzoyl peroxide. To achieve this, FlourSpec incorporates a novel end-to-end architecture that optimally balances spectral fidelity and classification performance while minimizing reconstruction errors. Additionally, a hybrid attention mechanism is employed to capture harmonic correlations across spectral bands, enhancing the system's ability to detect subtle adulteration signals. Comprehensive experimental evaluations demonstrate that FlourSpec achieves 97.44% accuracy in detecting multiple adulteration types and 93.12% accuracy in identifying various BPO concentrations. Notably, it is only 0.63% lower than expensive solutions and 18.74% higher than baseline systems at the same price, highlighting its effectiveness for trace-level detection. This performance underscores the effectiveness and robustness of FlourSpec in various environmental conditions, making it a valuable tool for ensuring flour quality and safety.

ACKNOWLEDGMENTS

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

REFERENCES

[1] M. Ačanski and D. Vujić. 2014. Comparing sugar components of cereal and pseudocereal flour by GC–MS analysis. *Food Chemistry* 145 (2014), 743–748.

[2] Jarmo T Alander, Vladimir Bochko, Birgitta Martinkauppi, Sirinnapa Saranwong, and Timo Mantere. 2013. A review of optical nondestructive visual and near-infrared methods for food quality and safety. *International Journal of Spectroscopy* 2013, 1 (2013), 341402.

[3] Boaz Arad and Ohad Ben-Shahar. 2016. Sparse recovery of hyperspectral signal from natural RGB images. In *Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part VII (Lecture Notes in Computer Science, Vol. 9911)*. Springer, 19–34.

[4] J. Bönick, G. Huschek, and H. Rawel. 2017. Determination of wheat, rye and spelt authenticity in bread by targeted peptide biomarkers. *Journal of Food Composition and Analysis* 58 (2017), 82–91.

[5] M. A. Brescia, A. Sgaramella, S. Ghelli, and A. Sacco. 2003. 1H HR-MAS NMR and isotopic investigation of bread and flour samples produced in southern Italy. *Journal of the Science of Food and Agriculture* 83 (2003), 1463–1468.

[6] Yuan-Yuan Chen and Zhi-Bin Wang. 2019. End-to-end quantitative analysis modeling of near-infrared spectroscopy based on convolutional neural network. *Journal of Chemometrics* 33, 5 (2019), e3122.

[7] M. Cocchi, C. Durante, G. Foca, A. Marchetti, L. Tassi, and A. Ulrici. 2006. Durum wheat adulteration detection by NIR spectroscopy multivariate calibration. *Talanta* 68 (2006), 1505–1511.

[8] European Commission. 2025. Glossary: Standard IARC classification. https://ec.europa.eu/health/scientific_committees/opinions_layman/en/electromagnetic-fields/glossary/ghi/iarc-classification.htm.

[9] G. Dinelli, A. Segura-Carretero, R. Di Silvestro, I. Marotti, D. Arráez-Román, S. Benedettelli, L. Ghiselli, and A. Fernadez-Gutierrez. 2011. Profiles of phenolic compounds in modern and old common wheat varieties determined by liquid chromatography coupled with time-of-flight mass spectrometry. *Journal of Chromatography A* 1218 (2011), 7670–7681.

[10] Karen Everstine, John Spink, and Shaun Kennedy. 2013. Economically motivated adulteration (EMA) of food: common characteristics of EMA incidents. *Journal of food protection* 76, 4 (2013), 723–735.

[11] Xian-guang Fan, Long Liu, Zhe-ming Kang, Ying-jie Zeng, Yu-liang Zhi, Ying-jie Xu, Jia-jie Zhang, and Xin Wang. 2020. Fast reconstruction of Raman spectra based on global weighted linear regression. *Chemometrics and Intelligent Laboratory Systems* 203 (2020), 104073.

[12] Market Data Forecast. 2025. Global Flour Market Size, Share, Trends & Growth Forecast Report Segmented By Raw Material (Wheat, Maize, Rice, Other Raw Materials), Application, Technology, And Region (North America, Europe, APAC, Latin America, Middle East And Africa), Industry Analysis From 2025 To 2033. https://www.marketdataforecast.com/market-reports/flour-market.

[13] Li Fu, Jiangwei Zhu, and Hassan Karimi-Maleh. 2021. An analytical method based on electrochemical sensor for the assessment of insect infestation in flour. *Biosensors* 11, 9 (2021), 325.

[14] Ying Fu, Yongrong Zheng, Lin Zhang, and Hua Huang. 2018. Spectral reflectance recovery from a single RGB image. *IEEE Transactions on Computational Imaging* 4, 3 (2018), 382–394.

[15] P. Geng, J. M. Harnly, and P. Chen. 2016. Differentiation of bread made with whole grain and refined wheat (T. aestivum) flour using LC/MS-based chromatographic fingerprinting and chemometric approaches. *Journal of Food Composition and Analysis* 47 (2016), 92–100.

[16] Xiao-Xi Guo, Wei Hu, Yuan Liu, Dong-Chen Gu, Su-Qin Sun, Chang-Hua Xu, and Xi-Chang Wang. 2015. Rapid analysis and quantification of fluorescent brighteners in wheat flour by Tri-step infrared spectroscopy and computer vision technology. *Journal of Molecular Structure* 1099 (2015), 393–398.

[17] Hong-Ju He, Yan Chen, Guanglei Li, Yuling Wang, Xingqi Ou, and Jingli Guo. 2023. Hyperspectral imaging combined with chemometrics for rapid detection of talcum powder adulterated in wheat flour. *Food Control* 144 (2023), 109378.

[18] M. Hernandez, T. Esteve, and M. Pla. 2005. Real-time polymerase chain reaction based assays for quantitative detection of barley, rice, sunflower, and wheat. *Journal of Agricultural and Food Chemistry* 53 (2005), 7003–7009.

[19] 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.

[20] 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.

[21] Yongzhi Huang, Kaixin Chen, Yandao Huang, Lu Wang, and Kaishun Wu. 2021. Vi-liquid: unknown liquid identification with your smartphone vibration. In *Proceedings of the 27th Annual International Conference on Mobile Computing and Networking (MobiCom '21)*. 174–187. https://doi.org/10.1145/3447993.3448621

[22] Neil F Johnson. 2021. Inhalation toxicity of talc. *Journal of aerosol medicine and pulmonary drug delivery* 34, 2 (2021), 79–107.

[23] R. Karoui, G. Cartaud, and E. Dufour. 2006. Front-face fluorescence spectroscopy as a rapid and nondestructive tool for differentiating various cereal products: a preliminary investigation. *Journal of Agricultural and Food Chemistry* 54 (2006), 2027–2034.

[24] Silvia Leccese. 2022. *Interaction between Orange Carotenoid Protein and mesoporous silica: from spectroscopic investigations to photoactive nanodevices*. Ph. D. Dissertation. Sorbonne Université.

[25] L. Lenhardt, I. Zekovic, T. Dramicanin, B. Milicevic, J. Burojevic, and M. Dramicanin. 2017. Characterization of cereal flours by fluorescence spectroscopy coupled with PARAFAC. *Food Chemistry* 229 (2017), 165–171.

[26] Yuqi Li, Chong Wang, and Jieyu Zhao. 2017. Locally linear embedded sparse coding for spectral reconstruction from RGB images. *IEEE Signal Processing Letters* 25, 3 (2017), 363–367.

[27] Debashis Nandi, Sumana Dutta, Shankar Prasad, Subhajyoti Ghosh, Aveer Mukherjee, Suranjana Maity, Dipankar Mandal, and Arpitam Chatterjee. 2025. An Approach of Flour Adulteration Detection in Turmeric Powder Using Computer Vision. In *International Conference on Systems and Technologies for Smart Agriculture*. 177–184. First Online: 29 January 2025.

[28] Mohammad Hossein Nargesi and Kamran Kheiralipour. 2024. Ability of visible imaging and machine learning in detection of chickpea flour adulterant in original cinnamon and pepper powders. *Research* 10, 16 (August 30 2024), e35944. Open access.

[29] AMS OSRAM. 2025. ams AS7421 Hyperspectral NIR Sensor. https://ams-osram.com/products/sensor-solutions/ambient-light-color-spectral-proximity-sensors/ams-as7421-hyperspectral-nir-sensor.

[30] K. Pastor, M. Ačanski, D. Vujić, G. Bekavac, S. Milovac, and S. Kravić. 2016. Rapid method for small grain and corn flour authentication using GC/EI–MS and multivariate analysis. *Food Analytical Methods* 9, 2 (2016), 443–450.

[31] Shibani Santurkar, Dimitris Tsipras, Andrew Ilyas, and Aleksander Madry. 2018. How does batch normalization help optimization? *Advances in neural information processing systems* 31 (2018).

[32] SCiO. 2025. SCiO Mini. https://shop.consumerphysics.com/.

[33] Yuanyuan Shao, Guantao Xuan, Zhichao Hu, and Yongxian Wang. 2019. Detection of adulterants and authenticity discrimination for coarse grain flours using NIR hyperspectral imaging. *Journal of Food Process Engineering* 42, 7 (2019), e13265.

[34] 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.

[35] Hui-Liang Shen, Pu-Qing Cai, Si-Jie Shao, and John H Xin. 2007. Reflectance reconstruction for multispectral imaging by adaptive Wiener estimation. *Optics express* 15, 23 (2007), 15545–15554.

[36] Zhan Shi, Chang Chen, Zhiwei Xiong, Dong Liu, and Feng Wu. 2018. HSCNN+: Advanced CNN-based hyperspectral recovery from RGB images. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*. 939–947.

[37] Rizwan Shoukat, Marta Cappai, Luca Pilia, and Giorgio Pia. 2025. Rice Starch Chemistry, Functional Properties, and Industrial Applications: A Review. *Polymers* 17, 1 (2025), 110.

[38] G. Sonnante, C. Montemurro, A. Morgese, W. Sabetta, A. Blanco, and A. Pasqualone. 2009. DNA microsatellite region for a reliable quantification of soft wheat adulteration in durum wheat-based foodstuffs by real-time PCR. *Journal of Agricultural and Food Chemistry* 57 (2009), 10199–10204.

[39] Wen-Hao Su and Da-Wen Sun. 2017. Evaluation of spectral imaging for inspection of adulterants in terms of common wheat flour, cassava flour and corn flour in organic Avatar wheat (Triticum spp.) flour. *Journal of Food Engineering* 200 (2017), 59–69.

[40] Bangjie Sun, Sean Rui Xiang Tan, Zhiwei Ren, Mun Choon Chan, and Jun Han. 2022. Detecting counterfeit liquid food products in a sealed bottle using a smartphone camera. In *Proceedings of the 20th Annual International Conference on Mobile Systems, Applications and Services (MobiSys '22)*. 42–55. https://doi.org/10.1145/3498361.3539776

[41] Yujie Tang, Yan He, Wenxing Yin, Yihong Gao, Xiangjun Zhang, and Wei Xia. 2022. Fabrication of a facile electrochemical sensor based on polyaniline/graphene oxide for sensitive detection of benzoyl peroxide in wheat flour. *Sensors and Actuators B: Chemical* 355 (2022), 130040. https://doi.org/10.1016/j.snb.2021.130040

[42] Yujie Tang, Yan He, Wenxing Yin, Yihong Gao, Xiangjun Zhang, and Wei Xia. 2023. Sensitive and selective detection of benzoyl peroxide in flour based on a dual-signal ratiometric electrochemical aptasensor. *Spectrochimica Acta Part A: Molecular and Biomolecular Spectroscopy* 293 (2023), 122089. https://doi.org/10.1016/j.saa.2023.122089

[43] Feifei Tao, Li Liu, Christopher Kucha, and Michael Ngadi. 2021. Rapid and non-destructive detection of cassava flour adulterants in wheat flour using a handheld MicroNIR spectrometer. *Biosystems Engineering* 203 (2021), 34–43.

[44] Johan Trygg and Svante Wold. 1998. PLS regression on wavelet compressed NIR spectra. *Chemometrics and Intelligent Laboratory Systems* 42, 1-2 (1998), 209–220.

[45] S. Verdú, F. Vásquez, R. Grau, E. Ivorra, A. J. Sánchez, and J. M. Barat. 2016. Detection of adulterations with different grains in wheat products based on the hyperspectral image technique: the specific cases of flour and bread. *Food Control* 62 (2016), 373–380.

[46] M. von Büren, M. Stadler, J. Lüthy, and P. Hubner. 2001. Detection of wheat adulteration of spelt flour and products by PCR. *European Food Research and Technology* 212 (2001), 234–239.

[47] Jonghyuk Yun, Kyoosik Lee, Kichang Lee, Bangjie Sun, Jaeho Jeon, Jeonggil Ko, Inseok Hwang, and Jun Han. 2024. PowDew: Detecting Counterfeit Powdered Food Products using a Commodity Smartphone. In *Proceedings of the 22nd Annual International Conference on Mobile Systems, Applications and Services*. 210–222.

[48] Wei-Feng Zhang and Dao-Qing Dai. 2008. Spectral reflectance estimation from camera responses by support vector regression and a composite model. *JOSA A* 25, 9 (2008), 2286–2296.

[49] Yuzhi Zhao, Lai-Man Po, Qiong Yan, Wei Liu, and Tingyu Lin. 2020. Hierarchical regression network for spectral reconstruction from RGB images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*. 422–423.

[50] Ling Zheng, Qian Bao, Shizhuang Weng, Jianpeng Tao, Dongyan Zhang, Linsheng Huang, and Jinling Zhao. 2022. Determination of adulteration in wheat flour using multi-grained cascade forest-related models coupled with the fusion information of hyperspectral imaging. *Spectrochimica Acta Part A: Molecular and Biomolecular Spectroscopy* 270 (2022), 120813.

[51] J. U. Ziegler, M. C. Leitenberger, F. H. Longin, T. Würschum, C. Reinhold, and R. 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.