Identifying the Defective: Detecting Damaged Grains for Cereal Appearance Inspection

Lei Fan ^{a,b,†;*}, Yiwen Ding^{a,†}, Dongdong Fan^a, Yong Wu^a, Maurice Pagnucco^b and Yang Song^b

^aGaozhe Technology, Hefei, China ^bThe School of Computer Science and Engineering, UNSW Sydney, Sydney, Australia [†]Equal Contribution. ORCiD ID: Lei Fan https://orcid.org/0000-0001-9472-7152

Abstract. Cereal grain plays a crucial role in the human diet as a major source of essential nutrients. Grain Appearance Inspection (GAI) serves as an essential process to determine grain quality and facilitate grain circulation and processing. However, GAI is routinely performed manually by inspectors with cumbersome procedures, which poses a significant bottleneck in smart agriculture.

In this paper, we endeavor to develop an automated GAI system: **AI4GrainInsp**. By analyzing the distinctive characteristics of grain kernels, we formulate GAI as a ubiquitous problem: Anomaly Detection (AD), in which healthy and edible kernels are considered *normal* samples while damaged grains or unknown objects are regarded as *anomalies*. We further propose an AD model, called AD-GAI, which is trained using only normal samples yet can identify anomalies during inference. Moreover, we customize a prototype device for data acquisition and create a large-scale dataset including 220*K* high-quality images of wheat and maize kernels. Through extensive experiments, AD-GAI achieves considerable performance in comparison with advanced AD methods, and **AI4GrainInsp** has highly consistent performance compared to human experts and excels at inspection efficiency over $20 \times$ speedup. The dataset, code and models will be released at **https://github.com/hellodfan/AI4GrainInsp**.

1 Introduction

Cereal grain plays a critical role in human survival and the development of civilizations, ensuring a reliable supply of food, contributing to poverty eradication and providing essential ingredients for various food products and daily necessities. The Quality Inspection of cereal Grains (OIG) is of paramount importance for standardizing grain storage, promoting fair circulation and guiding processing. It serves as a crucial metric for assessing nutrition, ensuring the security of supply, and identifying stratification (see Figure 1.a). Furthermore, OIG reflects crop conditions and holds the potential to guide sustainable and eco-friendly practices in smart agriculture. Currently, there are two dominant QIG methods: Chemical Analysis (CA) and Grain Appearance Inspection (GAI). CA is based on molecular biology and chemistry along with chemical substances and laboratory equipment, enabling highly precise inspection. On the other hand, GAI relies on visual characteristics to assess the appearance of grain kernels. Compared to CA, GAI is overwhelmingly adopted for high-throughput



Figure 1. a) Role of GAI in agriculture. b) GAI is formulated as an AD problem. Only *normal* samples are used for training, while models require identifying whether the test samples are *normal* or *anomalous* samples.

determination of the quality of cereal grains, including the detection of impurities, extraneous cereals, moldy grains and other damaged grains, as defined in the cereal ISO standard [14].

GAI is routinely executed manually by qualified inspectors. To illustrate this process, we consider the case of inspecting a shipment of raw wheat grains (originating from granaries or freighters). According to the sampling standard [13], the procedure involves taking a laboratory sample, which amounts to 60 grams and approximately 1600 kernels. These kernels are then inspected individually in a kernel-by-kernel procedure where inspectors must carefully examine the surface of kernels and then determine them as *healthy*, *damaged* or other categories (see Sec. 2.1). However, even qualified inspectors (with 5 to 10 years of expertise) typically require around 25 minutes to complete the inspection process, since the majority of grains are small in physical size, measuring less than $8 \times 4 \times 4mm^3$. As a result, inspecting these tiny grains demands a high level of concentration. Moreover, due to the nuances and superficial heterogeneity of cereal grains, manual inspection is prone to errors and lacks re-

^{*} Work done when interning at Gaozhe Technology. Corresponding Author. Email: lei.fan1@unsw.edu.au



Table 1. Wheat and maize examples of healthy, damaged grains and impurities (abbreviation used in the subsequent content).

liability. The available equipment or approaches for manual inspection are often cumbersome and limited in their capabilities. Therefore, in our work, we aim to develop automated GAI systems that can assist inspectors to enhance both the consistency and efficiency of inspections, providing significant social benefits.

Recently, Artificial Intelligence, particularly deep learning techniques [21, 10], has demonstrated an unprecedented level of proficiency, revolutionizing various fields such as medical image analysis [11, 9], autonomous driving [36] and anomaly detection in industries [1]. The widespread success of deep learning can primarily be attributed to the availability of large-scale high-quality datasets [33], sophisticated optimization objectives [23], and advanced model architectures [18, 6]. The application of deep learning techniques to GAI has the potential to significantly reduce labor costs and provide more stable and efficient decision-making compared to manual inspections. We thus aim to develop an automated GAI system equipped with deep learning techniques that can have the capability to replicate the decision-making of human experts. However, the challenge that how to acquire high-quality data hampers the development of automated GAI systems. The collection of data is critical in developing robust and accurate GAI systems. The data used to train these systems must be representative of the range of samples and environments that the system will encounter in the real world, and the data must be collected and labeled with great care to ensure that it is of high quality and sufficient quantity.

In this paper, we present an automated GAI system, named AI4GrainInsp. It consists of data acquisition using our custom-built device, data processing for dataset creation and a deep learning-based model for GAI. Specifically, we build an automated prototype device for data acquisition (see Figure 2), and further annotate a largescale dataset, called **OOD-GrainSet**, including 220K single-kernel wheat or maize images with object-centric masks and corresponding healthy or damaged category information. Moreover, by integrating the cross-domain knowledge between cereal science and deep learning techniques, we formulate GAI as a ubiquitous machine learning problem, Anomaly Detection (AD) [43], as shown in Figure 1.b. The objective of AD is to train a model using only normal samples, but the model is required to identify anomalous samples during inference. For GAI, the healthy and edible kernels are considered normal samples, while damaged kernels or other unknown objects are treated as anomalous samples.

We further propose an AD model for GAI, called **AD-GAI**, with a customized data augmentation strategy to synthesize anomaly-like samples based on normal samples from both image-level and featurelevel perspectives. These synthesized data are used as negative samples for training a discriminator in a supervised manner. We conduct extensive experiments to verify the superior performance of AD-GAI on our **OOD-GrainSet** and the publicly available MVTec AD datasets that is typically used as the benchmark dataset for AD. **AI4GrainInsp** shows strong potential both in consistency and efficiency in comparison with human experts. The main contributions are listed as follows:

- We propose an automated GAI system: **AI4GrainInsp**, which is a complete pipeline from data acquisition to deep learning-based data analysis models.
- We formulate GAI as an AD problem and further propose a data augmentation-based model for GAI, called AD-GAI. Extensive experiments are conducted to verify the superiority of AD-GAI on both our grain dataset and a public benchmark dataset for AD, and validate the feasibility and efficacy of AI4GrainInsp both in consistency and efficiency.
- We release a large-scale dataset, called *OOD-GrainSet*, including 220K images for wheat and maize with expert-level annotations.

2 Background

2.1 Grain Appearance Inspection

Wheat and maize are two of the main cereal grains and together make up approximately 42.5% of the world's crop yield in 2022 reported in [12]. GAI serves as a requisite procedure [14] for ensuring grain quality, requiring inspectors to inspect the surface of grains carefully and classify them into healthy, damaged grains, impurities and other contaminants. Damaged grains refer to grains of decreased value and can be mainly categorized into six types: sprouted (SD) grain, fusarium & shriveled grain, black point (BP) grain for wheat or heated (HD) grain for maize, moldy (MY) grain, broken (BN) grain, grain attacked by pests (AP), as illustrated in Table 1. F&S, MY and BP grains are contaminated by fusarium or fungus, while SD, HD, BN and AP grains have decreased values in various nutrients. On the other hand, impurities (IM), including organic objects (foreign cereals) and unknown objects (stone, plastic), can also have deleterious effects on grain processing and circulation. Similar to healthy grains, BN, AP, BP and HD are also edible to some extent. Therefore, we conduct two data partition schedules in experiments, *i.e.*, healthy grains vs. damaged grains, and edible grains vs. inedible grains.

In this paper, we propose an automated system, **AI4GrainInsp**, that utilizes a sampling device coupled with deep learning techniques. Considering the heterogeneity and diversity of grains, we formulate GAI as an anomaly detection problem, and demonstrate our **AI4GrainInsp** equipped with deep learning techniques increases the accuracy and efficiency of the inspection process.

2.2 AI for Smart Agriculture and Food

In recent years, artificial intelligence (AI) techniques have achieved significant progress in the field of smart agriculture [28]. For example, by analyzing satellite images or drone images, AI can forecast and monitor climatic and soil conditions, as well as predict crop yield and production [7]. AI-based Unmanned Aerial Vehicles (UAV) and autonomous tractors have provided robust navigation and dynamic planning techniques for smart irrigation and disease control [41]. With the help of remote cameras, AI is also capable of analyzing plant diseases and detecting pest distributions [20]. Furthermore, some researchers have attempted to use AI to recognize food categories [27], and estimate the calorie and nutrition content [39].

However, there has been limited research [8] on cereal grains in the cultivation-grain-processed food streamline. Grain quality determination still lags behind, with no automated devices currently available and manual-inspection strategies proving to be cumbersome. In this paper, we focus on this critical yet underestimated field of grain quality determination, especially GAI. We demonstrate that building an automated GAI system is a highly challenging problem. We endeavor to build an effective system powered by deep learning techniques, to ensure food safety and contribute to the development of smart agriculture and promoting progress toward "Good Health and Well-being" and Sustainable Development Goals.

2.3 Anomaly Detection

Visual anomaly detection [43, 24, 16] means that only normal samples are available during training time, while normal and anomalous samples should be identified during inference. Early studies attempt to formulate anomaly detection as one-class classification [32, 30, 26] that assign high confidence to *in-distribution* samples and low probabilities to *out-of-distribution* samples, and there is a line of work called SVDD-based methods [38, 32] that train models to project representations into a hypersphere space.

The majority of recent deep learning-based studies adopt reconstruction-based methods. These methods are built on a hypothesis that models can effectively estimate the distribution of normal samples. These methods [17, 45, 46, 5] typically adopt an encoderdecoder architecture (*e.g.*, autoencoder) to encode and decode normal images and low-dimension representations sequentially. To better learn representations, some studies [17, 31] introduce a memory mechanism to explicitly store different patterns of anomaly-free samples. Similar to reconstruction-based methods, some studies [2, 34] tried to learn and localize discrepancies between normal and anomalous samples by relying on knowledge distillation [15].

Recently, data augmentation-based strategy has also been widely explored [22, 45, 25, 48]. These methods try to synthesize anomalylike samples based on normal samples by using well-designed data augmentation techniques, and these synthesized samples are used as supervision signals to train classification models. For example, Cut-Paste [22] employs CutMix [44], DRAEM [45] and DeSTSeg [48] generates anomaly-like samples based adding noise on normal images. SimpleNet [25] tries to identify normal features extracted from normal samples or anomalous features generated by adding noises to normal features. In this paper, our proposed AD-GAI tries to synthe-



Figure 2. The blueprint and prototype device for data acquisition.

size anomaly-like samples from both image-level and feature-level perspectives, achieving considerable performance on three datasets.

3 AI4GrainInsp

Al4GrainInsp consists of three components: a prototype device, a large-scale dataset, **OOD-GrainSet**, and an AD grain analysis framework, **AD-GAI**. We first introduce our prototype device for data acquisition. Using this device, we captured and annotated a large-scale dataset containing about 220K images of single kernels with expert-level annotations. We then describe our proposed AD-GAI for automated grain quality determination.

3.1 Data Acquisition

There are two main challenges for capturing the visual information of raw grains: capturing high-quality images and collecting digital images efficiently. To overcome these challenges, we developed a customized prototype device for digitizing the surface information of grains, as shown in Figure 2. For the first challenge, we employ a dual-camera strategy where two industrial cameras (860 DPI) with corresponding light sources are vertically placed along a transparent plate. We refer to the two cameras as the UP and DOWN cameras. To tackle the second challenge, we employ a conveyor belt with vibration bands. This enables the transparent plate to maintain a horizontal loop movement between the ends of the conveyor belt and the two cameras. The vibration bands can effectively separate the grain kernels and force kernels onto the transparent plate individually. As a result, a batch of grain kernels in the plate can be digitized together at a high sampling rate for data acquisition.

Taking a laboratory wheat sample as an example (about $60 \pm 0.5g$ and near 1600 kernels), to digitize the images of kernels, we divided them into several batches using the conveyor belt. Each batch consists of about 150 to 300 kernels delivered onto the transparent plate by the conveyor belt with vibration bands. Then, the plate piled with wheat kernels is placed at the center of the dual cameras. Each camera with the corresponding light source is controlled to capture high-quality images of grain kernels in a large receptive field, producing a pair of UP (I_{up}) and DOWN (I_{down}) images from the two cameras for a batch of kernels. Finally, we obtain several pairs of high-quality images for a laboratory wheat sample.

3.2 OOD-GrainSet

Raw Data: Figure 3 illustrates an example of a pair of images captured from UP (I_{up}) and DOWN (I_{down}) angles, each of which has a high resolution of 3644×5480 pixels covering $91 \times 137 mm^2$. As



Figure 3. Annotation examples for a batch of grain kernels. Both high-resolution UP (I_{up}) and DOWN (I_{down}) images contain a set of grain kernels. The pair information, object localization and morphological shape are healthy or damaged grains (denoted in different colors) are provided.

 I_{up} and I_{down} are totally vertical to the transparent plate, the combination of UP and DOWN images covers about 92 to 98% of the superficial areas of grain kernels according to physical measurements.

Expert-level Annotations: As a pair of UP and DOWN images $(I_{up} \text{ and } I_{down})$ capture the surface information from the top and bottom views, each kernel in these images has two sides and shares the same healthy or damaged grain category information. Inspired by rotation object detection [42] and instance segmentation [23] tasks, we annotate these images from four perspectives: single-kernel pair information, object localization, kernel mask and damaged grain category (see Figure 3). For example, a pair of UP and DOWN images produce a set of single-kernel images containing two sides of kernels in a horizontal layer, where each image has a corresponding segmentation mask M depicting the morphological shape at the pixel-level. All single kernels are classified as healthy, impurities, or one of the six damaged grain categories. To simplify the processing for building the AD dataset, all single-kernel images are processed with geometrical transformations to show similar poses, as shown in Table 1.

Distributions of OOD-GrainSet. Our dataset¹, called *OOD-GrainSet*, involves two types of cereal grains: wheat and maize. Given the nature of grains, the proportion of damaged grains is relatively small and we made efforts to maintain a balanced distribution for building OOD-GrainSet, as shown in Figure 4. For wheat data, we annotated about 180K single-kernel images, including 145K healthy grains and 5K images for each damaged grain category and impurities. For maize data, we annotated about 40K single-kernel images, including 33K healthy grains and 1K images for each damaged grain category and impurities. Moreover, we additionally annotated several wheat and maize samples that are used for **AI4GrainInsp** vs. Experts experiments (see Sec. 4.4).

3.3 AD-GAI

Different from public anomaly detection data [1] containing rich color and contextual information collected from the wild, we consider that *normal* (healthy or edible grains) and *anomalous* (damaged grains or unknown objects) samples in **OOD-GrainSet** share mostly common visual information in terms of shape and context. The primary distinctions between normal and anomalous samples are fine-grained, and the characteristics of damaged grains (such as F&S or AP) are subtle in size, such as wormholes or moldy spots.



Figure 4. The distributions of *OOD-GrainSet*, including healthy, damaged grains (DG) and impurities for wheat and maize. Among damaged grains, the categories of BN, AP, and BP/HD are classified as edible.

Based on such analysis and *a priori* understanding, we propose AD-GAI which is based on a data augmentation-based strategy to synthesize anomaly-like samples. Inspired by previous AD methods for industrial inspection [45, 22, 25, 48], as there are no anomalous samples that can be used during the training phase, we employ the data augmentation strategy to synthesize anomaly-like samples based on normal samples from both image-level and feature-level perspectives. These synthesized anomaly-like samples and normal samples are used together as supervision signals for training a classification model in an end-to-end manner, after which the model can identify whether test samples are normal or anomalous during inference.

Notation. We denote $\mathbb{D} = \{I_0, \ldots, I_{N-1}\}$ as a training set containing only N normal images and a test set $\mathbb{T} = \{I_0, \ldots, I_{M-1}\}$ containing M normal or anomalous images. Each image $I \in \mathbb{R}^{W \times H \times C}$ (W, H and C of width, height and channels respectively) has a label $y \in \{0, 1\}$ where 0 and 1 mean normal or anomalous. Our goal is to train a model using only \mathbb{D} during training, while the model can classify whether test samples in \mathbb{T} are normal or anomalous.

The overview of AD-GAI and pseudo-code of the training procedure are presented in Figure 5 and Algorithm 1 respectively. During the training phase, given an input image I_x with the label y = 0, I_x is augmented by adding noise to synthesize an anomaly-like image I_n . Both I_x and I_n are fed into the feature extractor ϕ_{ex} to extract patchaware features \mathcal{F}_x and \mathcal{F}_n respectively. Then, \mathcal{F}_x is augmented by adding Gaussian noise to synthesize an anomaly-like feature \mathcal{F}_a . Finally, these features \mathcal{F}_x , \mathcal{F}_n and \mathcal{F}_a are concatenated and then fed into the classifier ϕ_{cls} , which is optimized to discriminate these features as normal or anomalous. The details of these methods are described in the following.

Simulation of image-level anomalies. We follow the previous methods [45, 48] to synthesize anomaly-like image I_n based on a normal image I_x , as shown in Figure 6. A binary mask $M_b \in \mathbb{R}^{W \times H}$ generated from Perlin noise P contains several anomaly shapes, and

Algorithm 1 Pseudo-code of AD-GAI in a PyTorch-like style.
<pre>E, C: the feature extractor and classifier .params = ResNet50.params # initialize E with pre- trained ResNets for I_x in data_loader: # load input normal images I_a = aug_add_noise(I_x) # obtain anomaly-like images : I_a</pre>
<pre>F_x = E.forward(I_x) # patch-aware features: F_x F_n = E.forward(I_a) # patch-aware features: F_n F_a = F_x + random(N) # obtain anomaly-like features: F_a</pre>
<pre>loss = CrossEntropyLoss(C.forward([F_x,F_n,F_a])). mean() # channel-wise CE loss loss.backward() # update C and E</pre>

¹ More details can be found on the project website.



Figure 5. Overview of AD-GAI. The input normal image I_x is augmented to synthesize an anomaly-like image I_n . Both I_x and I_n are fed into the feature extractor ϕ_{ex} to obtain patch-aware features \mathcal{F}_x and \mathcal{F}_n respectively. Then, \mathcal{F}_x is further augmented by adding Gaussian noise to synthesize an anomaly-like feature \mathcal{F}_a . Finally, the classifier ϕ_{cls} , consists of MLP layers, is trained to predict \mathcal{F}_x as *negative*, \mathcal{F}_n and \mathcal{F}_a as *positive*.

an arbitrary image I_a sampled from the another dataset (*e.g.*, ImageNet [33]) is blended with I_x based on M_b , which is defined as:

$$I_n = (1 - M'_b) \odot I_x + \beta (M'_b \odot I_a) + (1 - \beta) (M'_b \odot I_x), \quad (1$$

where β is the opacity parameter between [0.15, 1] as described in [45, 47], and \odot is the element-wise multiplication operation. M'_b is generated by conducting pixel-wise *and* operation between M_b and the mask M (provided in annotations) of the grain, which limits that generated anomaly shapes fall onto the foreground of grains.

Extraction of patch-aware features. We follow prior methods [31, 25, 4] and use a pre-trained model (*e.g.*, ResNet50 [18] trained on ImageNet) to extract features for images. Specifically, the feature extractor ϕ_{ex} employs a ResNet-like model to extract hierarchical features $\{f_l \in \mathbb{R}^{w_l \times h_l \times c_l}, l \in (1, 2, 3, 4)\}$ truncated from different convolutional stages. These features are further aggregated to obtain patch-aware features with larger receptive-of-views. For a position (i, j) with the entry $f_l^{i,j}$, the aggregation is defined as:

$$f'_{l} = \phi_{avg}(\{f^{(i,j)}_{l} | (i,j) \in \mathbf{p}^{*}\}),$$
(2)

where ϕ_{avg} denotes the adaptive average pooling, and \mathbf{p}^* denotes a patch centered at (i, j) and a size of p (set to 3). The aggregation retains the resolutions of features. To enrich feature information, we fuse aggregated features from different stages to obtain final patch-aware features. For simplification, we leverage l = 3, 4 features that contain abundant spatial and semantic information. The high-level features (l = 4) with the smaller resolution are interpolated to the same resolution of low-level features (l = 3), which is defined as:

$$\mathcal{F}_x = concat[f'_l, \phi_{int}(f'_{l+1})], \tag{3}$$

where ϕ_{int} and *concat* denote the linear interpolation and channel concatenation operation respectively. The patch-aware features \mathcal{F}_n can also be extracted for the synthesized anomaly images I_n .

Simulation of feature-level anomalies. Inspired by previous methods [25, 4], we attempt to synthesize anomaly-like samples from the feature perspective. For the patch-aware features \mathcal{F}_x extracted from the normal image I_x , the anomaly-like features \mathcal{F}_a are synthesized by adding noise on \mathcal{F}_x , which is defined as:

$$\mathcal{F}_{a}^{(i,j)} = \mathcal{F}_{x}^{(i,j)} + \epsilon, \tag{4}$$

where ϵ is sampled from Gaussian distribution $\mathcal{N}(\mu, \sigma^2)$. We visualize the similarities among normal image features \mathcal{F}_x , anomaly-like image features \mathcal{F}_n and synthesized anomaly-like features \mathcal{F}_a , and



Figure 6. The simulation of image-level anomalies. M_b is a binary mask used for indicating the blend between the input image I_x (with a mask M) and an arbitrary image I_a to synthesize an anomaly-like image I_n .

these features are extracted by using t-SNE technique [40], as shown in Figure 8.b.

Optimization objective. These features are finally fed into the classifier ϕ_{cls} . The classifier ϕ_{cls} employs a multi-layer perceptron (MLP) layer, and is trained to predict *negative* for normal features \mathcal{F}_x and *positive* for anomaly-like image features \mathcal{F}_n and anomaly-like features \mathcal{F}_a . We empirically employ cross-entropy loss (CE) as the optimization objective:

$$L = \frac{1}{w_l \cdot h_l} \sum_{(i,j)}^{w_l, h_l} CE(\phi_{cls}(\mathcal{F}^{(i,j,:)})).$$
(5)

where $\mathcal{F}^{(i,j,:)}$ denotes the feature vector at position (i, j), and (w_l, h_l) is the spatial shape of features.

Inference. During inference, the branches of simulation of imagelevel and feature-level anomalies are discarded. The output of $\max \phi_{cls} \mathcal{F}_{I_t}^{(i,j)} | (i,j) \in \mathcal{P}^*$ is used as the anomaly score for the test sample I_t where \mathcal{P}^* is the set of positions of \mathcal{F} .

4 Experiments

4.1 Experimental Settings

Datasets. We explore our AD-GAI on our OOD-GrainSet and the public benchmark MVTec AD [1]. The **MVTec AD dataset** [1] is widely used for evaluating anomaly detection methods. It provides

Methods	Туре	MVTec AD	OOD-GrainSet				
in the and the second			Wheat(set1)	Wheat(set2)	Maize(set1)	Maize(set2)	Average
Deep-SVDD (ICML-18 [32])	Distance-based	59.2	86.5	86.6	80.9	76.4	82.6
PADiM (ICPR-21 [4])	Distance-based	95.8	73.1	67.5	67.2	59.4	66.8
Mem-AE (ICCV-19 [17])	Reconstruction-based	-	85.8	84.9	73.8	56.4	75.2
DRAEM (ICCV-21 [45])	Reconstruction-based	98.1	79.8	59.5	66.4	78.1	70.9
RevDist (CVPR-22 [5])	Reconstruction-based	98.4	90.1	89.2	86.5	81.9	86.9
CSI (NeurIPS-20 [37])	Data Augmentation-based	-	83.6	77.3	84.7	78.6	81.1
CutPaste (CVPR-21 [22])	Data Augmentation-based	96.1	76.7	77.5	75.1	71.3	75.2
AD-GAI (single R50 model)	Data Augmentation-based	99.0	94.2	93.5	87.7	82.5	89.5
AD-GAI (ensemble R50&R101)	Data Augmentation-based	<u>99.1</u>	<u>95.9</u>	<u>94.1</u>	<u>88.2</u>	<u>82.8</u>	<u>90.2</u>

Table 2. Comparisons of our AD-GAI and advanced methods on MVTec AD dataset and four sub-sets of our OOD-GrainSet.

5354 high-resolution images across 10 object categories and 5 texture categories, such as toothbrush and wood. The training set comprises 3629 normal images, while the test set includes 1725 normal or anomalous images along with pixel-level anomaly annotations. We follow the experimental settings as [45, 5] where we train an individual model for each category.

For our **OOD-GrainSet**, we construct four sub-sets according to the grains' conditions, as shown in Table 3. Wheat(set1) and Maize(set1) indicate that only healthy grains are treated as *nor-mal* samples, while the remaining grains are considered anomalies. Wheat(set2) and Maize(set2) mean that healthy grains and some of the damaged yet edible grains are combined as *normal* samples. We split normal samples into 70% and 30% partitions for the training and test sets, preserving the ratio of different categories. All anomalous samples are used for the test set.

Implementation details. All experiments are conducted on a workstation with RTX 3090 GPUs based on the PyTorch platform [29]. We use ResNet-50 [18] pre-trained on ImageNet [33] as the default feature extractor. We employ Adam optimizer [19] with momentum of (0.8, 0.999), weight decay of 1×10^{-4} , initial learning rate of 1×10^{-3} . The batch size is set to 4 and the training epoch is set to 8 and 16 for wheat and maize respectively.

Evaluation metrics. we use the commonly used Area Under the Receiver Operating Curve (AUROC) as the metric for both MVTec AD and OOD-GrainSet. In addition, to validate **AI4GrainInsp** with human experts, we employ the Macro F1-score (threshold is set to 0.3 for our AD-GAI) as the metric, and we also report the inspection time to evaluate the runtime efficiency.

Table 3. Detailed settings of two subsets: **set1** and **set2**. \checkmark and \bigcirc indicate that the category is used as *normal* or *anomalous* samples. Only healthy grains are treated as *normal* samples in the **set1**, while edible grains are treated as *normal* samples in the **set2**.

Dataset	HY Damaged Grains							IM
Settings		BN	AP	BP/HD	SD	FS	MY	
Wheat/Maize(set1) Wheat/Maize(set2)	\checkmark	O ✓	O ✓	$\overset{\bigcirc}{\checkmark}$	00	00	00	00

4.2 Comparisons with Advanced Methods

The experiments were conducted on MVTec AD and four sub-sets of OOD-GrainSet, by comparing with three types of AD methods: distance-based (Deep-SVDD [32] and PADiM [4]), reconstruction-based (Mem-AE [17], DRAEM [45] and RevDist [5]) and data augmentation-based methods (CSI [37] and CutPaste [22]).

As shown in Table 2, our **AD-GAI** produces the best performance on all four sub-sets of OOD-GrainSet, achieving about 5.8%, 4.9%, 1.7% and 0.9% improvement over other advanced methods

on Wheat(set1), Wheat(set2), Maize(set1) and Maize(set2) respectively. Moreover, our AD-GAI also produces excellent results on the MVTec AD dataset, with 99.1% of image-level AUROC performance. Compared to other data augmentation-based methods, our model achieves substantial improvements, which validates the effectiveness of our AD-GAI that attempts to synthesize anomaly-like samples from both image-level and feature-level perspectives.

4.3 Ablation Study

Backbones for feature extractors. The feature extractor ϕ_{ex} extracts patch-aware features from input images. We test different ResNet-like [18] backbones without data augmentation, as shown in Table 4. We observe that using R50 pre-trained on ImageNet gains significant improvements of 14.4% and 19.5% on Wheat(set1) and Maize(set1) compared to R50 from scratch, which confirms the effectiveness of using pre-trained models to extract features. We further explore backbones with different parameter scales. Using lightweight model R18 pre-trained on ImageNet also outperforms R50 from scratch, and larger models R50 and R101 can produce better performance than R18. It is noted that R50 and R101 show similar performance, and we select R50 as our default backbone due to its relatively lower computational costs.

 Table 4.
 Ablation study of different backbones and data augmentation techniques on Wheat(set1)/Maize(set1).

Backbone	R50 from scratch	R18	R50	R101
	79.8/67.2	88.9/80.7	94.2/87.7	94.1/86.3
Data	None	Flip+Rot	Mixup [47]	RandAug [3]
Augmentation	94.2/87.7	90.3/87.4	93.5/87.1	89.7/83.2

Data augmentations. We conducted experiments by using different data augmentation techniques. Compared to Flip+Rot (horizontal and vertical flipping and 90, 180, 270 rotations), mixup [47] or RandAug [3], it is noted that using no data augmentation (in addition to our sample synthesis) produces the best performance of 94.2% and 87.7% on Wheat(set1) and Maize(set1) respectively. We consider it is because both training and test samples are well-processed during data annotations, and using heavy data augmentation techniques can be harmful to simulations of anomaly-like samples since the distinctions between normal and anomalous samples are subtle.

Structure of AD-GAI and noise levels. We also conducted experiments on Wheat(set1) to investigate the impact of noise levels on both image-level and feature-level simulations of anomalies, as shown in Figure 7. We formulate a noise parameter r that represents the ratio of the maximum area of binary noise mask M_b to the mask of grain M. Particularly, using only image-level (*i.e.*, $\sigma = 0$) or feature-level (*i.e.*, r = 0) simulations produces moderate results



Figure 7. The ablation study of noise levels on Wheat(set1). r and σ can enable or control noise levels of image-level and feature-level simulations.

of 89.7% and 90.2% respectively, which confirms the effectiveness of using both simulations together. AD-GAI achieves the best performance of 94.2% when $\sigma = 0.025$ and r = 0.2. We consider that small values of σ or r cannot synthesize anomalies well, while large values will produce redundant anomaly-like samples harmful to training an effective classifier with limited normal samples.

Qualitative analysis. We utilize the Grad-CAM technique [35] to visualize anomalous samples and prediction results from two datasets, as shown in Figure 8.a. Our AD-GAI effectively focuses on discriminative regions, such as wormholes, moldy points, scratches, etc. Moreover, we employ the t-SNE technique [40] to qualitatively demonstrate the similarities among features from simulations and real anomalous samples, as shown in Figure 8.b. We observe that both image-level and feature-level synthesized anomalies are closer to real anomalous samples than normal samples, which verifies the effectiveness of our data augmentation strategies.



Figure 8. a) Visualization of anomaly images, prediction of AD-GAI with Grad-CAM technique [35] and experts' annotations. b) Visualization of features from normal, anomalous and simulations by using t-SNE [40].

4.4 AI4GrainInsp versus. Human Experts

We further evaluate our **AI4GrainInsp** system in comparison with human experts. We enlisted two junior inspectors, JI1 and JI2, who had 3 years of experience, and two senior inspectors, SI1 and SI2, who had near 10 years of experience. We built two prototype devices, D1 and D2, each equipped with deployed AD-GAI models. We collected 4 groups of wheat samples (each of 60g) with 4%, 8%, 15% and 20% proportions of damaged grains and impurities, and 3 groups of maize samples (each of 600g, since maize grains are heavier) with 4%, 8% and 15% proportions of damaged grains and impurities. We conducted and averaged two individual inspections for each test sample. We report the F1-score and time cost, which is the total running time of the system or inspectors.

As shown in Figure 9, our **AI4GrainInsp** produces impressive performance, which is highly consistent with senior inspectors SI1 and SI2 while being much more time-efficient over $20 \times$ speedup (about 73s vs. 1550s). Similar to wheat, experimental results on maize also validate the superiority and efficiency of our system. In contrast, the results of two junior inspectors JI1 and JI2 are relatively moderate with fluctuation, and their inspection time costs are much higher than those of devices. Therefore, we consider that our system has the potential to assist inspectors in grain quality determinations.



Figure 9. AI4GrainInsp vs. human experts on wheat and maize grains.

5 Conclusions and Future Work

In this paper, we propose a comprehensive automated GAI system called **AI4GrainInsp**, which includes a prototype device for data acquisition, a high-quality dataset for evaluation, and an anomaly detection model **AD-GAI**. Our model utilizes data augmentation techniques to synthesize anomaly-like samples by adding noise to normal samples from both image-level and feature-level perspectives. Experimental results demonstrate the superiority of AD-GAI, and **AI4GrainInsp** is highly consistent with human experts with much higher efficiency. Additionally, we release a large-scale dataset, **OOD-GrainSet**, containing 220K single-kernel images across eight categories for wheat and maize.

There still exist many challenges in GAI. For example, our **AI4GrainInsp** system coupled with a customized device has high manufacturing costs, and we aim to develop low-cost solutions, such as using smartphones to enable widespread deployment. Moreover, we also plan to train and apply **AI4GrainInsp** to more types of cereal grains, such as rice, sorghum, etc. The key challenge is to collect abundant grains from different geographical locations and build a comprehensive cereal grain atlas. We hope that our work will draw more attention to GAI-related fields and promote smart agriculture, contributing to reaching the SDG goals.

References

- Paul Bergmann, Michael Fauser, David Sattlegger, and Carsten Steger, 'MVTec AD-a comprehensive real-world dataset for unsupervised anomaly detection', in *CVPR*, pp. 9592–9600, (2019).
- [2] Paul Bergmann, Michael Fauser, David Sattlegger, and Carsten Steger, 'Uninformed students: Student-teacher anomaly detection with discriminative latent embeddings', in *CVPR*, pp. 4183–4192, (2020).
- [3] Ekin D Cubuk, Barret Zoph, Jonathon Shlens, and Quoc V Le, 'Randaugment: Practical automated data augmentation with a reduced search space', in *CVPR Workshops*, pp. 702–703, (2020).
- [4] Thomas Defard, Aleksandr Setkov, Angelique Loesch, and Romaric Audigier, 'Padim: a patch distribution modeling framework for anomaly detection and localization', in *ICPR*, pp. 475–489. Springer, (2021).
- [5] Hanqiu Deng and Xingyu Li, 'Anomaly detection via reverse distillation from one-class embedding', in CVPR, pp. 9737–9746, (2022).
- [6] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, et al., 'An image is worth 16x16 words: Transformers for image recognition at scale', in *ICLR*, (2020).
- [7] Ngozi Clara Eli-Chukwu, 'Applications of artificial intelligence in agriculture: A review', *Engineering, Technology & Applied Science Re*search, 9(4), 4377–4383, (2019).
- [8] Lei Fan, Yiwen Ding, Dongdong Fan, Donglin Di, Maurice Pagnucco, and Yang Song, 'Grainspace: A large-scale dataset for fine-grained and domain-adaptive recognition of cereal grains', in *CVPR*, pp. 21116– 21125, (2022).
- [9] Lei Fan, Arcot Sowmya, Erik Meijering, and Yang Song, 'Learning visual features by colorization for slide-consistent survival prediction from whole slide images', in *MICCAI*, pp. 592–601. Springer, (2021).
- [10] Lei Fan, Arcot Sowmya, Erik Meijering, and Yang Song, 'Fast ff-toffpe whole slide image translation via laplacian pyramid and contrastive learning', in *MICCAI*, pp. 409–419. Springer, (2022).
- [11] Lei Fan, Arcot Sowmya, Erik Meijering, and Yang Song, 'Cancer survival prediction from whole slide images with self-supervised learning and slide consistency', *IEEE Transactions on Medical Imaging*, 42(5), 1401–1412, (2023).
- [12] Food and Agriculture Organization. World food situation, https://www. fao.org/worldfoodsituation/csdb/en. Accessed on 2023-05-10.
- [13] International Organization for Standardization, 'ISO 24333: Cereals and cereal products — Sampling', Standard, International Organization for Standardization, (December 2009).
- [14] International Organization for Standardization, 'ISO 5527: Cereals Vocabulary', Standard, International Organization for Standardization, (February 2015).
- [15] Hinton Geoffrey, Vinyals Oriol, and Jeffrey Dean, 'Distilling the knowledge in a neural network', in NIPS Deep Learning and Representation Learning Workshop, (2015.).
- [16] Mariana-Iuliana Georgescu, Antonio Barbalau, Radu Tudor Ionescu, et al., 'Anomaly detection in video via self-supervised and multi-task learning', in *CVPR*, pp. 12742–12752, (2021).
- [17] Dong Gong, Lingqiao Liu, Vuong Le, et al., 'Memorizing normality to detect anomaly: Memory-augmented deep autoencoder for unsupervised anomaly detection', in *ICCV*, pp. 1705–1714, (2019).
- [18] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun, 'Deep residual learning for image recognition', in CVPR, pp. 770–778, (2016).
- [19] Diederik P Kingma and Jimmy Ba, 'Adam: A method for stochastic optimization', arXiv preprint arXiv:1412.6980, (2014).
- [20] Cardim F. L., Maria Elisa Damascena de Almeida Leandro, Valero, et al., 'Automatic detection and monitoring of insect pests—A review', *Agriculture*, **10**(5), 161, (2020).
- [21] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton, 'Deep learning', nature, 521(7553), 436–444, (2015).
- [22] Chun-Liang Li, Kihyuk Sohn, Jinsung Yoon, and Tomas Pfister, 'Cutpaste: Self-supervised learning for anomaly detection and localization', in *CVPR*, pp. 9664–9674, (2021).
- [23] Tsung-Yi Lin, Michael Maire, Serge Belongie, et al., 'Microsoft COCO: Common objects in context', in ECCV, pp. 740–755. Springer, (2014).
- [24] Wen Liu, Weixin Luo, Dongze Lian, and Shenghua Gao, 'Future frame prediction for anomaly detection – a new baseline', in CVPR, (2018).
- [25] Zhikang Liu, Yiming Zhou, Yuansheng Xu, and Zilei Wang, 'Simplenet: A simple network for image anomaly detection and localization', in *CVPR*, pp. 20402–20411, (2023).

- [26] Philipp Liznerski, Lukas Ruff, Robert A Vandermeulen, et al., 'Explainable deep one-class classification', in *ICLR*, (2020).
- [27] Weiqing Min, Zhiling Wang, Yuxin Liu, et al., 'Large scale visual food recognition', *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(8), 9932–9949, (2023).
- [28] A. Mitra, Sukrutha LT Vangipuram, et al., 'Everything you wanted to know about smart agriculture', arXiv preprint arXiv:2201.04754, (2022).
- [29] Adam Paszke, Sam Gross, Massa, et al., 'Pytorch: An imperative style, high-performance deep learning library', *NeurIPS*, 32, (2019).
- [30] Pramuditha Perera, Ramesh Nallapati, and Bing Xiang, 'OCGAN: Oneclass novelty detection using GANs with constrained latent representations', in *CVPR*, pp. 2898–2906, (2019).
- [31] Karsten Roth, Latha Pemula, Joaquin Zepeda, Bernhard Schölkopf, Thomas Brox, and Peter Gehler, 'Towards total recall in industrial anomaly detection', in *CVPR*, pp. 14318–14328, (2022).
- [32] Lukas Ruff, Robert Vandermeulen, Nico Goernitz, Lucas Deecke, Shoaib Ahmed Siddiqui, Alexander Binder, Emmanuel Müller, and Marius Kloft, 'Deep one-class classification', in *ICML*, pp. 4393–4402. PMLR, (2018).
- [33] Olga Russakovsky, Jia Deng, Hao Su, et al., 'Imagenet large scale visual recognition challenge', *IJCV*, 115(3), 211–252, (2015).
- [34] Mohammadreza Salehi, Niousha Sadjadi, Soroosh Baselizadeh, Mohammad H Rohban, and Hamid R Rabiee, 'Multiresolution knowledge distillation for anomaly detection', in *CVPR*, pp. 14902–14912, (2021).
- [35] Ramprasaath R Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra, 'Grad-cam: Visual explanations from deep networks via gradient-based localization', in *ICCV*, pp. 618–626, (2017).
- [36] Pei Sun, Henrik Kretzschmar, Xerxes Dotiwalla, et al., 'Scalability in perception for autonomous driving: Waymo open dataset', in *CVPR*, pp. 2446–2454, (2020).
- [37] Jihoon Tack, Sangwoo Mo, Jongheon Jeong, and Jinwoo Shin, 'CSI: Novelty detection via contrastive learning on distributionally shifted instances', *NeurIPS*, 33, 11839–11852, (2020).
- [38] David MJ Tax and Robert PW Duin, 'Support vector data description', Machine learning, 54(1), 45–66, (2004).
- [39] Quin Thames, Arjun Karpur, Wade Norris, et al., 'Nutrition5k: Towards automatic nutritional understanding of generic food', in *CVPR*, pp. 8903–8911, (2021).
- [40] Laurens Van der Maaten and Geoffrey Hinton, 'Visualizing data using t-SNE.', JMLR, 9(11), (2008).
- [41] Ahmad Latif Virk, Mehmood Ali Noor, Sajid Fiaz, et al., 'Smart farming: an overview', Smart Village Technology, 191–201, (2020).
- [42] Gui-Song Xia, Xiang Bai, Jian Ding, et al., 'DOTA: A large-scale dataset for object detection in aerial images', in *CVPR*, pp. 3974–3983, (2018).
- [43] Jingkang Yang, Kaiyang Zhou, Yixuan Li, and Ziwei Liu, 'Generalized out-of-distribution detection: A survey', arXiv preprint arXiv:2110.11334, (2021).
- [44] Sangdoo Yun, Dongyoon Han, Seong Joon Oh, Sanghyuk Chun, Junsuk Choe, and Youngjoon Yoo, 'Cutmix: Regularization strategy to train strong classifiers with localizable features', in *ICCV*, pp. 6023–6032, (2019).
- [45] Vitjan Zavrtanik, Matej Kristan, and Danijel Skočaj, 'DRAEM-A discriminatively trained reconstruction embedding for surface anomaly detection', in *ICCV*, pp. 8330–8339, (2021).
- [46] Vitjan Zavrtanik, Matej Kristan, and Danijel Skočaj, 'Reconstruction by inpainting for visual anomaly detection', *Pattern Recognition*, **112**, 107706, (2021).
- [47] Hongyi Zhang, Moustapha Cisse, Yann N Dauphin, and David Lopez-Paz, 'Mixup: Beyond empirical risk minimization', in *ICLR*, (2018).
- [48] Xuan Zhang, Shiyu Li, Xi Li, Ping Huang, Jiulong Shan, and Ting Chen, 'Destseg: Segmentation guided denoising student-teacher for anomaly detection', arXiv preprint arXiv:2211.11317, (2022).