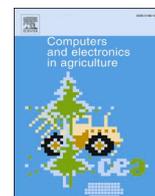




Contents lists available at ScienceDirect

Computers and Electronics in Agriculture

journal homepage: www.elsevier.com/locate/compag

Generative adversarial networks (GANs) for image augmentation in agriculture: A systematic review

Yuzhen Lu ^{a,*}, Dong Chen ^{b,1}, Ebenezer Olaniyi ^{a,1}, Yanbo Huang ^c^a Department of Agricultural and Biological Engineering, Mississippi State University, Mississippi State, MS 39762, USA^b Department of Electrical and Computer Engineering, Michigan State University, MI 48824, USA^c United States Department of Agriculture Agricultural Research Service, Genetics and Sustainable Agriculture Research Unit, Mississippi State, MS 39762, USA

ARTICLE INFO

Keywords:

GAN
Image Augmentation
Agriculture
Deep Learning
Computer Vision

ABSTRACT

In agricultural image analysis, optimal model performance is keenly pursued for better fulfilling visual recognition tasks (e.g., image classification, segmentation, object detection and localization), in the presence of challenges with biological variability and unstructured environments. Large-scale, balanced and ground-truthed image datasets are tremendously beneficial but most often difficult to obtain to fuel the development of highly performant models. As artificial intelligence through deep learning is impacting analysis and modeling of agricultural images, image augmentation plays a crucial role in boosting model performance while reducing manual efforts for image collection and labelling, by algorithmically creating and expanding datasets. Beyond traditional data augmentation techniques, generative adversarial network (GAN) invented in 2014 in the computer vision community, provides a suite of novel approaches that can learn good data representations and generate highly realistic samples. Since 2017, there has been a growth of research into GANs for image augmentation or synthesis in agriculture for improved model performance. This paper presents an overview of the evolution of GAN architectures followed by a first systematic review of various applications in agriculture and food systems (<https://github.com/Derekabc/GANs-Agriculture>), involving a diversity of visual recognition tasks for plant health conditions, weeds, fruits (preharvest), aquaculture, animal farming, plant phenotyping as well as postharvest detection of fruit defects. Challenges and opportunities of GANs are discussed for future research.

1. Introduction

Machine vision or imaging technologies are ubiquitous in agricultural applications. A diversity of imaging techniques, e.g., panchromatic/monochromatic, color or RGB (red-green-blue), near-infrared (NIR) multispectral/hyperspectral imaging, are available to replace human vision to acquire information from agricultural processes and food products (Chen et al., 2002; Davies, 2009; Mavridou et al., 2019; Lu et al., 2020); they can deliver improved efficacy, efficiency, and objectivity in vision tasks in agriculture, such as crop stress diagnosis, weed control, phenotyping, robotic fruit harvesting, and grading and sorting of agricultural produce. Integrated with machine vision technologies, machine learning offers a suite of methods/algorithms (Liakos et al., 2018; Rahman et al., 2019) for analyzing and learning patterns from data and has been widely utilized in data-driven visual recognition (e.g.,

image classification, object detection/localization, and segmentation) and quantification tasks (e.g., plant stress rating/quantification) in agriculture-food systems. Traditional machine learning methodologies for image analysis require the extraction of hand-engineered features (e.g., color, shape, texture) from images to build pattern recognition and classification models (O'Mahony et al., 2019), which is a cumbersome, trial-and-error process without guaranteeing optimal model performance.

In the past ten years, with the advancements in computing hardware including GPUs, artificial intelligence (AI) with deep learning (DL) (Bengio et al., 2021; LeCun et al., 2015; Krizhevsky et al., 2012) has emerged as a new paradigm and workhorse for pattern classification and learning from data. Popularized by supporting frameworks such as PyTorch (Paszke et al., 2019) and TensorFlow (Abadi et al., 2016), DL is advantageous as it enables seamless integration of feature extraction and pattern classification processes while learning high-quality image

* Corresponding author at: Department of Agricultural and Biological Engineering, Mississippi State University, Mississippi State, MS 39762, USA.

E-mail address: y1747@msstate.edu (Y. Lu).¹ Equal contribution.<https://doi.org/10.1016/j.compag.2022.107208>

Received 14 April 2022; Received in revised form 28 June 2022; Accepted 7 July 2022

Available online 2 August 2022

0168-1699/© 2022 Elsevier B.V. All rights reserved.

Nomenclature	
AC-GAN	Auxiliary classifier generative adversarial network
AdaIN	Adaptive instance normalization
ADA	Adversarial domain adaptation
AI	Artificial intelligence
AP	Average precision
AR-GAN	Attentive recurrent generative adversarial network
BN	Batch normalization
CGAN	Conditional generative adversarial network
CNN	Convolutional neural network
CoGAN	Coupled generative adversarial network
DCGAN	Deep convolutional generative adversarial network
DL	Deep learning
ESR-GAN	Enhanced super-resolution generative adversarial network
FCN	Fully connected network
FID	Fréchet inception distance
GAN	Generative adversarial network
InfoGAN	Information maximizing generative adversarial network
IoU	Intersection over union
IS	Inception score
mAP	Mean average precision
mIoU	Mean intersection over union
NIR	Near-infrared
PCA	Principal component analysis
PNAS	Progressive neural architecture search
ProGAN	Progressively growing generative adversarial network
ReLU	Rectified linear unit
RGB	Red-Green-Blue
SA-GAN	Self-attention generative adversarial network
SN	Spectral normalization
SR-GAN	Super-resolution generative adversarial network
SSIM	Structural similarity index
StyleGAN	Style generative adversarial network
SVM	Support vector machine
TransGAN	Transformer generative adversarial network
UAV	Unmanned aerial vehicle
WGAN	Wasserstein generative adversarial network
WANG-GP	Wasserstein generative adversarial network with gradient penalty

representations and achieving state-of-the-art performance in visual recognition tasks (Krizhevsky et al., 2012; He et al., 2016; Carion et al., 2020; Zhai et al., 2021), which otherwise could be challenging using traditional machine learning methods. The impacts of DL, especially convolutional neural networks (CNNs), on tackling computer vision tasks in agriculture, have been well reviewed in Kamilari and Prenafeta-Boldú (2018), Koirala et al. (2019) and Zhang et al. (2020). Despite impressive results in analyzing and modeling agricultural images, it has been clearly demonstrated that supplies of large-scale datasets are necessary for ensuring the performance of DL (Sun et al., 2017) or advanced machine learning models (Paullada et al., 2021) while avoiding overfitting. To achieve human-level performance, as a rough rule of thumb, thousands of images per category are expected to train DL models from scratch (LeCun et al., 2015; Goodfellow et al., 2016). Given the positive impact of data size, a number of immensely large image datasets with millions of labeled images have been created in the computer vision community, e.g., ImageNet (Deng et al., 2009), COCO (Lin et al., 2014), LVIS (Gupta et al., 2019), Open Images (Kuznetsova et al., 2020), just to name a few.

Analysis and modeling of agricultural images face a myriad of challenges resulting from significant biological variability (no two leaves are the same) and unstructured environments (e.g., object occlusion, variable lighting conditions, cluttered scenes) (Bechar and Vigneault, 2016; Barth et al., 2020; Vougioukas, 2019). To address these challenges for building robust models accentuates the need for large-scale datasets capturing sufficient variations under a wide range of imaging conditions. Collecting and annotating such datasets with ground-truths, however, are time-consuming, resource-intensive, and costly. This is especially true for specific applications (e.g., plant disease detection, weed recognition, fruit defect detection), which place constraints on biological materials and imaging conditions, and for precise annotations at pixel level. Currently, there are very few image datasets that consist of hundreds if not thousands of images per category, on par with the aforementioned datasets in computer vision, and are publicly available (Lu and Young, 2020). The scarcity of annotated, large-scale image datasets relevant to agricultural processes is posing a crucial bottleneck to harnessing the power of advanced AI methods including DL algorithms in the agriculture and food domain. One common approach to mitigating the insufficiency of physically collected data is data/image augmentation that algorithmically expands the scale and variations of datasets (Lu and Young, 2020). A suite of image

augmentation techniques (Simard et al., 2003; Wong et al., 2016; Shorten and Khoshgoftaar, 2019; Khalifa et al., 2021) has been proposed for improving the performance of DL models.

Collectively, there are two categories of image augmentation techniques, including basic or traditional image augmentation and the augmentation based on DL algorithms. The former involves various image transformations using geometric or color processing approaches (Shorten and Khoshgoftaar, 2019). The second relies on DL techniques (Khalifa et al., 2021), among which generative adversarial networks (GANs) offer a novel method for image augmentation through generative model learning of underlying distributions of training data. In supervised machine learning/DL pipelines, the original dataset is initially partitioned into training, validation, and testing sets (or training and testing sets). Image augmentation is normally conducted for training (and or validation) images during model training and optimization (LeCun et al., 2015), with test data unaugmented to avoid data leakage, although image augmentation is found to be useful at test time for assessing model performance (Shanmugam et al., 2022).

Originally proposed by Goodfellow et al. (2014), GAN is a new framework of generative modeling (Tomczak, 2022), which aims to synthesize new data with the same characteristics of training instances (usually images), visually resembling the data in the training set. Since its advent, GAN has enjoyed great popularity for generating realistic data, and various GAN-based methods have been proposed for image syntheses in the past five years (Creswell et al., 2018; Huang et al., 2018; Yi et al., 2019; Cui et al., 2021), with applications spreading rapidly from computer vision and machine learning communities to domain-specific areas such as biomedical diagnosis (Yi et al., 2019; Bissoto et al., 2021), bioinformatics (Lan et al., 2020) as well as agriculture. While traditional image augmentation techniques are extensively used in model development (Simard et al., 2003; Ciresan et al., 2010), the diversity or variations garnered through simple geometric/color transformations can be limited with little additional information. Instead, GANs provide novel and promising ways through representation learning to potentially induce more variations and enrich the original datasets, benefiting downstream modeling tasks (Perez and Wang, 2017; Bowles et al., 2018; Frid-Adar et al., 2018). While a few articles describe theoretical aspects and generic applications of GANs (LeCun et al., 2015; Goodfellow et al., 2020; Huang et al., 2018; Gui et al., 2020; Creswell et al., 2018; Yi et al., 2019; Liu et al., 2021), none of them, to the best of our knowledge, are specifically devoted to the application of GANs for

image augmentation in the agricultural and food domain.

This paper is therefore to provide the first systematic review of GANs for image augmentation/synthesis in agriculture and food systems to facilitate DL-assisted computer/machine vision tasks. We survey recent developments of GAN architectures and their application in precision agriculture, phenotyping, and postharvest handling, and discuss technical challenges facing GANs and future research needs. This review would be beneficial for the research communities to exploit GAN techniques to address challenges posed by the shortage of labeled, large-scale image datasets.

The remainder of the paper is organized as follows. Section 2 describes the method and exclusion criteria for the literature review. Section 3 describes basic image augmentation methods for machine learning systems. Described next is a comprehensive overview of representative GAN architectures in Section 4. A systematic review of agricultural applications of GANs is presented in Section 5, followed by a discussion of challenges and further work of GANs in Section 6 and a conclusive summary in Section 7.

2. Methods

This review adopts the PRISMA (preferred reporting items for systematic reviews and meta-analysis) guideline in Moher et al. (2009), an evidence-based framework to perform a systematic review. Fig. 1 shows the PRISMA guideline flowchart. All the authors agreed upon the design and method that follow the guideline for this review. Likewise, all the authors agreed on the focused questions, search databases, and inclusion/exclusion criteria for article selection.

The focused questions considered in this review include: can new training images be generated to expand the training data? What impacts of image augmentation on DL for agricultural applications? How can GANs help to expand training data for improving DL model performance? Regarding information sources, the literature search was conducted following the PRISMA guideline (Moher et al., 2009). We initially searched several major scientific databases (e.g., Web of Science, ScienceDirect, Google Scholar, Springer) to collate the literature. We discovered that some databases cover limited articles since GAN in agriculture is currently an emerging application area. Articles covered in one database with limited articles might be also covered by another database with broader coverage. Therefore, we discarded the databases with limited and repeated articles. Finally, three mainstream scientific journal databases including Google Scholar, ScienceDirect and IEEE Xplore, along with the open access arXiv database, were considered in

this review. Relevant articles were searched and chosen using the search strategy and inclusion criteria described next.

To collect relevant literature, we searched the databases based on selected keywords, which were entered into the search engine of each database. We first did searching with individual keywords such as “synthetic images of plant diseases”, “weed detection data augmentation”, “synthetic images of weeds”, “plant disease data augmentation”, “data augmentation in agriculture”, “generative adversarial network in agriculture”, etc. To collect more publications, we also grouped similar keywords with AND and the keywords in diverse groups with OR. This resulted in search string combinations such as “synthetic images of weeds” AND “artificial images in weed control”, “generative adversarial network in pest recognition” OR “synthetic images in plant seedlings”, “artificial images in postharvest” OR “aquaculture synthetic images”, “image augmentation in plant seedling” OR “artificial images in weed control”, “generative adversarial network in fruit detection” AND “synthetic images in fruit detection”, to name a few.

We restricted our search to only articles written in English and published within five past years (2017–2022) due to the recent emergence of GANs in the agricultural field. Finally, we collected 303 articles from the database search, comprising 139 articles from Google Scholar, 50 articles from ScienceDirect, 81 articles from IEEE Xplore and 33 articles from arXiv. Articles screening was then carried out based on duplicate, title, abstract, and full text of the papers using the following exclusion criteria: (1) review articles, (2) articles irrelevant to data augmentation in agriculture, (3) publications that are duplicate or already enlisted from another search database, and (4) publications that do not include the full text.

After the eligibility and exclusion criteria were thoroughly implemented, the literature screening was left with a total number of 59 research articles at the time of writing (Figs. 1 and 2). The articles reviewed in this study are summarized on the accompanying GitHub repository (<https://github.com/Derekabc/GANs-Agriculture>), which is actively maintained.

3. Basic image augmentation

Basic image augmentation is loosely referred to as transformations that manipulate the geometry and color characteristics of images (Taylor and Nitschke, 2018; Khalifa et al., 2021), corresponding to geometric and photometric transformations, respectively. The geometric transformations (also known as image warping) involve altering the geometry and sometimes associated spatial information (e.g., masks,

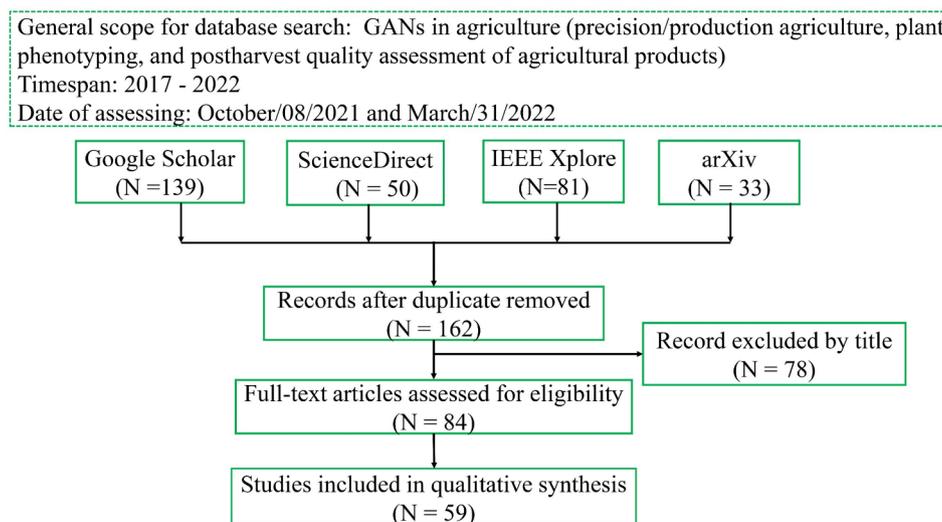


Fig. 1. The PRISMA guideline flowchart used in this review. The figure first row illustrates initial selected articles based on the keywords that enhanced the initial filtering before other exclusion criteria are applied.

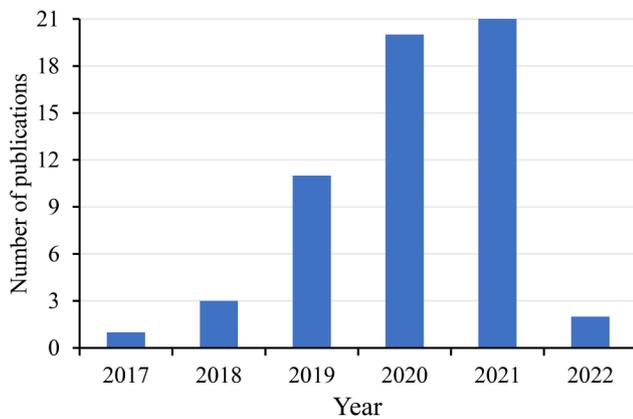


Fig. 2. Numbers of publications over the years.

bounding boxes) of an image. The object shape in the image may be preserved or not, depending on the linearity of transformation methods. Affine transformations, which preserve spatial collinearity and are traditionally used in image registration and mosaicking (Shapiro and Stockman, 2001), are most often used for image augmentation when training deep learning models, including translation, rotation, magnification, and the composition of any combination and sequence. There are also non-linear warping methods (Oliveira et al., 2017), e.g., elastic deformations in Simard et al. (2003). The photometric methods involve the transformations across color channels of an RGB image (Taylor and Nitschke, 2018), or spectral channels for multi-/hyper-spectral images. These methods do not change object geometry and spatial characteristics. Color transformations are commonly done through color jittering that is to randomly adjust color saturation, brightness, and contrast in image color spaces (Affi et al., 2019; Cubuk et al., 2019), color channel swapping, principal component analysis (PCA) color augmentation (also known as Fancy PCA) (Krizhevsky et al., 2012; Taylor and Nitschke, 2018), among others. Common image processing methods, such as filtering, histogram equalization, edge enhancement, noise addition, etc., which have the effect of introducing variations to image characteristics (e.g., shape, color and or texture), can also be used for image augmentation.

Fig. 3 illustrates common basic image augmentation performed on two weed images, randomly chosen in CottonWeed dataset (Chen et al., 2022), for which the augmented images can be used for weeds classification and detection tasks. The augmentation was achieved using

Albumentation (Buslaev et al., 2020), an open-source library for image augmentation. Although DL frameworks, such as TensorFlow and PyTorch, have the functionality of basic image augmentation, a number of standalone, software packages dedicated to image augmentation tasks using a rich variety of spatial/pixel-level transformation operations, have been developed and publicly available to enhance image augmentation. Table 1 summarizes common open-source packages/tools for image augmentation.

While basic augmentation methods described above can produce significant amounts of data and help train DL models, these methods still have their drawbacks. The basic image transformations (e.g., translation, cropping, rotation) tend to generate highly correlated samples and may not account for variations resulting from different imaging protocols or sequences, which are thus not guaranteed to be effective for model improvements or generalization, if improperly implemented, and even can lead to overfitting (Shorten and Khoshgoftaar, 2019). Generally, these augmentation techniques are applied on one image at a time and thus inadequate to learn the variations or invariant features from the rest of training data. Sometimes, heavy augmentations may eliminate or damage meaningful semantic contents, which is undesirable for images acquired under strict standards (Konidakis et al., 2019; Yi et al., 2019). Little new information would be gained from otherwise small modifications to the images (e.g., translation of an image by a few pixels, image cropping).

Generative modeling, an important branch of machine learning that aims to learn the underlying distribution (probability) from data and

Table 1

Example of open-source software packages for image augmentation.

Packages/Tools	URL	Reference
Albumentations	https://github.com/albumentations-team	Buslaev et al. (2020)
Augmentor	https://github.com/mdbloice/Augmentor	Bloice et al. (2019)
CutBlue	https://github.com/clovaai/cutblur	Yoo et al. (2020)
Gryds	https://github.com/tueimage/gryds	Eppenhof & Plum (2019)
imgaug	https://github.com/aleju/imgaug	NA
Pymia	https://github.com/rundherum/pymia	Jungo et al. (2021)
Rising	https://github.com/PhoenixDL/rising	NA
TorchIO	https://github.com/fepegar/torchio	Perez-Garcia et al. (2021)

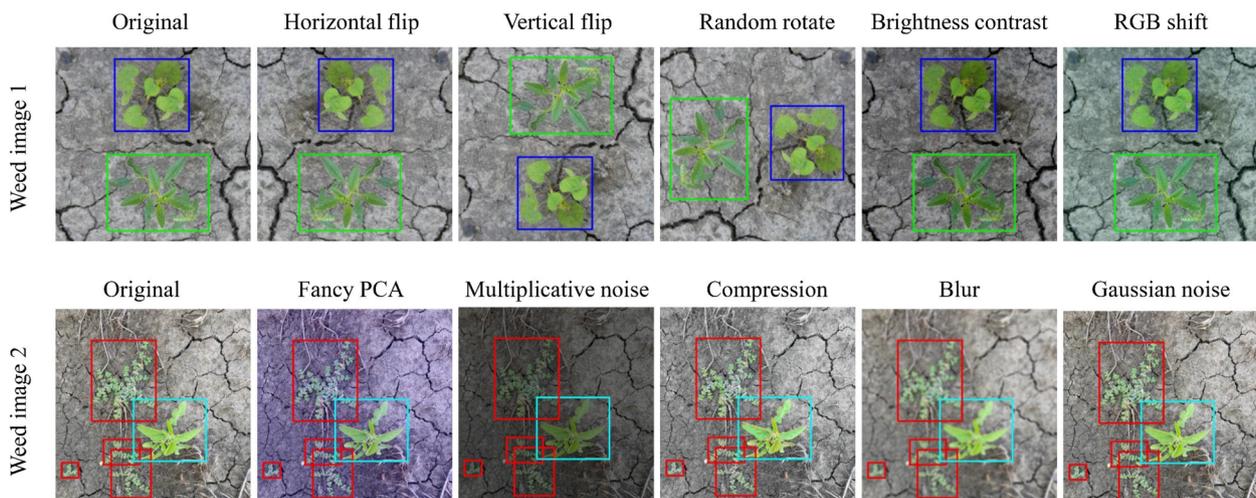


Fig. 3. Illustration of basic image augmentation on two weed images, each of which is processed by five different traditional image augmentation methods. The bounding boxes of difficult colors in the images indicate weed plants of different weed species, and the bounding box information is preserved during image augmentation.

generate new samples, has achieved striking progress in recent years and offered alternative approaches to image generation in computer vision applications. There is a great variety of generative modeling techniques (Tomczak, 2022), many of which focus on a *log-likelihood* function for generative learning, such as normalizing flows, autoregressive models, variational autoencoders, energy-based models, and lately, diffusion probabilistic models (Ho et al., 2020). Generative adversarial networks (GANs) represent a distinct group of generative modeling techniques, which utilize adversarial training with an *adversarial* loss to assess model and data distributions (Goodfellow et al., 2014). Compared to other generative models, GANs, which are described in the following sections, are arguably among the most successfully used for synthesizing realistic-looking images and promise to be a valuable addition to the arsenal of basic color and geometric transformation techniques for image augmentation.

4. Principle and architectures of GANs

This section describes the basic concepts and principles of genuine GAN and important architectural developments. Readers already familiar with theoretic aspects of GANs may skip to the next section. As an emerging class of unsupervised modeling techniques, GANs (Goodfellow et al., 2014; Goodfellow et al., 2020) has brought breakthroughs to data generation for deep learning. The vanilla GAN (Goodfellow et al., 2014) consists of two machine learning models (generally neural networks) called *generator* (G) and *discriminator* (D), which are trained in an adversarial process. As illustrated in Fig. 4, G is fed with random noise (z) and outputs synthetic data, while D fed with real samples (x) is to discriminate between the real and fake samples, $G(z)$, generated from the G . Simply put, G is trained to trick the D , while the D is optimized not to be fooled by the G .

The original way to train a GAN model is to find a discriminator with the maximized classification and a generator that maximally confuses the discriminator. The training process involves optimizing the following value function (Goodfellow et al., 2014):

$$V(D, G) = E_{x \sim p_x} [\log D(x)] + E_{z \sim p_z} [\log(1 - D(G(z)))] \quad (1)$$

where p_x and p_z denote the real data and generated data distributions, respectively. The D is trained to maximize the probability of assigning the correct labels to fake samples from the G as well as training samples, and simultaneously G is trained to minimize the loss $\log(1 - D(G(z)))$, which leads to the two-player minimax game (Goodfellow et al., 2014). During training, G and D models are updated iteratively, with the parameters of one model updated and the parameters of the other fixed. When the two models G and D are fully trained, the game reaches a Nash equilibrium (Mirza & Osindero, 2014; Cao et al., 2018) where D is unable to distinguish the samples from the real data distribution or generated distribution from G , thus always predicting 0.5 for all samples.

The networks D and G are often parameterized by differentiable functions which can be fully connected, convolutional or recurrent networks. The original GAN (Goodfellow et al., 2014) used fully connected networks as its building block, which are only suitable for

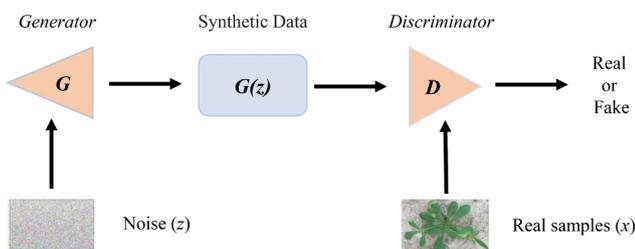


Fig. 4. The framework of a vanilla generative adversarial network (GAN) that consists of two main models, i.e., the generator (G) and the discriminator (D).

handling low resolution images datasets, e.g., MNIST and CIFAR-10, since fully connected networks have limited image representation capacity. To synthesize high-resolution, complex images, convolutional neural networks (CNNs) that have more powerful representation ability have become core components in recent GAN models, providing better performance in image generation. Extended from the vanilla GAN model (Goodfellow et al., 2014), a myriad of GAN architecture variants has been proposed in literature for different applications (Wang et al., 2020; Huang et al., 2018; Creswell et al., 2018). Presented below are several important GAN architectures that have been used to generate synthetic images in agricultural applications as well as other disciplines (e.g., computer vision, medical imaging, biology) to improve model performance.

4.1. Conditional generative adversarial network

The original GAN has no control over the data being generated. By conditioning the GAN on auxiliary information, the model can be enhanced to generate images with desired attributes. Mirza and Osindero (2014) who are among the first to develop conditional GANs (CGANs), conditioned both the G and D models on class label y , as illustrated in Fig. 5. The G receives the input of random noise z and the class label y concatenated in joint representation, and the D takes real samples and their corresponding class label. The objective function of CGAN is defined as follows:

$$L(D, G, y) = E_{x \sim p_x} [\log D(x|y)] + E_{z \sim p_z} [\log(1 - D(G(z|y)))] \quad (2)$$

By Mirza and Osindero (2014), the class-conditional GAN was able to generate MNIST digits with the class labels encoded as one-hot vectors and it was capable of learning multi-modal representations using user generated multi-label prediction. Auxiliary classifier GAN (or AC-GAN) modifies CGAN by predicting class labels with the discriminator instead of taking them as input (Odena et al., 2017), which could generate discriminable and diverse ImageNet samples exhibiting global coherence. In addition to class labels, the conditional variable y can be images (Isola et al., 2017; Zhu et al., 2017), bounding boxes, key points, or image text descriptions. Reed et al. (2016) generated plausible images (64×64 pixels) conditioned on text descriptions. Zhang et al. (2017) proposed StackGANs for text-to-image synthesis, which could generate images of $4 \times$ better resolution. Isola et al. (2017) proposed the first general-purpose image conditional GAN framework, called Pix2Pix (<https://github.com/phillipi/pix2pix>), as illustrated in Fig. 6, for image-to-image translation (translating representations from source to output images). Pix2Pix requires using a training set of aligned image pairs for learning mappings between input and output images. Subsequently, Zhu et al. (2017) eliminated the requirement to learn mappings using unpaired training data (Section 4.4). Image-conditional GANs have been widely applied for image synthesis in face editing, image inpainting and super resolution (Huang et al., 2018; Gui et al., 2020).

4.2. Deep convolutional generative adversarial network

Using convolutional layers in place of fully connected layers in the original GAN leads to the development of a family of GAN architectures known as deep convolutional GAN (DCGAN) (Radford et al., 2015), one

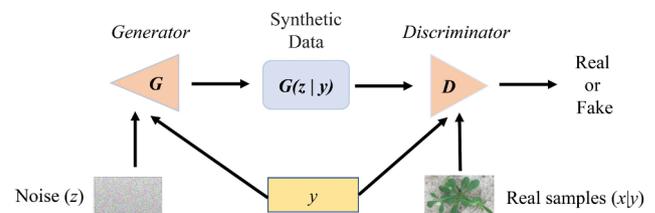


Fig. 5. The framework of conditional generative adversarial network with both the generator and discriminator conditioned on class labels (y).

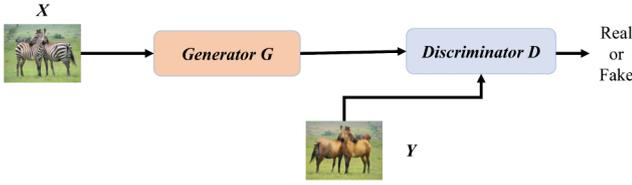


Fig. 6. The framework of Pix2Pix (Isola et al., 2017) for image-to-image translation where aligned or paired training data are required for training.

of the most successful GAN variants widely used by many later models. In the DCGAN, convolutional layers are adopted to represent D and G to learn unsupervised representations, and specifically strided (D) and fractionally-strided (G) convolutions (Shelhamer et al., 2016) were utilized to learn up-sampling and down-sampling operations, which may contribute to improvements in the quality of synthesized images. To avoid model collapse (all input images mapped to the same output image, a common challenge with training GANs), in the DCGAN, batch normalization (BN) (Ioffe and Szegedy, 2015) was applied in both D and G during training, helping stabilize the GAN training and preventing collapsing all samples to a single sample. Radford et al. (2015) also demonstrated that proper activation functions in G [rectifying linear unit (ReLU activation)] and D (leaky rectified activation) can help models to learn quickly. Tested on various large-scale image datasets, DCGAN was found effective for learning the hierarchy of representations from object parts to scenes (Radford et al., 2015).

4.3. Coupled generative adversarial network

Extending the original GAN model, coupled generative adversarial network (CoGAN) (Liu and Tuzel, 2016) aims to learn a joint distribution of multi-domain images and generate tuples of images in different domains without correspondence supervision. The architecture of CoGAN for translating images in two domains, as illustrated in Fig. 7, consists of a pair of GANs, i.e., GAN1 and GAN2, each responsible for image synthesis in one domain. During training, the two generators ($G1$ and $G2$) form a team for generating a pair of images in two different domains, while the two discriminators ($D1$ and $D2$) are to differentiate the images in the respective domain from those produced by the corresponding generator. Identical network structure and weights are shared in the first layers of the generators and in the last few layers of the discriminators so as to enforce CoGAN to learn a joint distribution of images even without correspondence supervision between the generators and discriminators. The two-domain CoGAN can be generalized to multiple domains by stacking more GAN modules and enforcing the weight-sharing constraint. The CoGAN showed good performance in several joint distribution learning tasks (e.g., color and depth images) as well as image translation and domain adaptation applications (Liu and Tuzel, 2016).

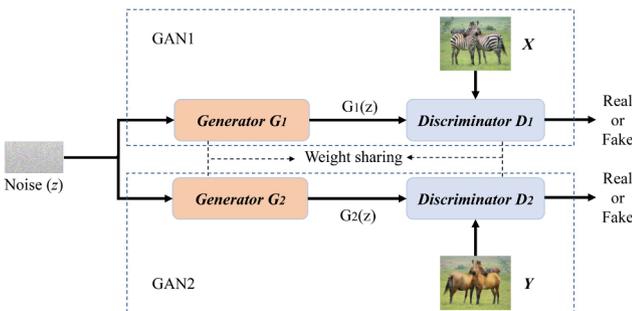


Fig. 7. The framework of coupled generative adversarial network (CoGAN) for translating two image domains.

4.4. Cycle-consistent generative adversarial network

Built upon Pix2Pix (Isola et al., 2017), cycle-consistent generative adversarial network (CycleGAN) was designed for image-to-image translation with unpaired training data (Zhu et al., 2017). CycleGAN, which is trained in an unsupervised manner, eliminates the need for paired training data in Pix2Pix and is particularly useful in applications where paired training data are not readily available. CycleGAN, as shown in Fig. 8, consists of two generator networks, G_{XY} and G_{YX} , and two discriminators, D_X and D_Y , where X and Y are source and target domains, respectively. The key idea of CycleGAN is that the two mapping functions, $G_{XY}: X \rightarrow Y$ and $G_{YX}: Y \rightarrow X$ should be inverse of each other, and the translation is to be cycle consistent (in a sense that, if we translate, e.g., zebras to horses and then translate horses back to zebras, we should get the exact same original zebra images). The adversarial loss (Goodfellow et al., 2014) is applied to both mapping functions, leading to the objectives as follows:

$$L(D_X, G_{YX}) = E_{x \sim p_x} [\log D_X(x)] + E_{y \sim p_y} [\log(1 - D_X(G_{YX}(y)))] \quad (3)$$

$$L(D_Y, G_{XY}) = E_{x \sim p_x} [\log D_Y(y)] + E_{y \sim p_y} [\log(1 - D_Y(G_{XY}(x)))] \quad (4)$$

To achieve cycle-consistent mappings, two L1 reconstruction losses (also named cycle consistency losses) are introduced to the optimization (Zhu et al., 2017), with the objective function defined as follows:

$$L(G_{XY}, G_{YX}) = E_{x \sim p_x} [\|G_{YX}(G_{XY}(x) - x)\|_{L_1}] + E_{y \sim p_y} [\|G_{XY}(G_{YX}(y) - y)\|_{L_1}] \quad (5)$$

Thus, the overall objective is expressed as below:

$$L_{\text{CycleGAN}} = L(D_X, G_{YX}) + L(D_Y, G_{XY}) + \beta L(G_{XY}, G_{YX}) \quad (6)$$

where β is the weighting coefficient balancing the relative importance of objective functions.

The CycleGAN has achieved impressive performance in image synthesis and style transfer (Zhu et al., 2017). Following the idea of dual learning, CycleGAN shares structural similarities with DualGAN (Yi et al., 2017) for unpaired image-to-image translation, except that the former uses sigmoid cross-entropy loss format while the latter uses the loss advocated by the WGAN (Wasserstein GAN) (Arjovsky and Bottou, 2017). The original CycleGAN that is unconditional has been extended to a conditional variant for image generation (Lu et al., 2018).

4.5. Information maximizing generative adversarial networks

Information maximizing generative adversarial networks (InfoGAN), proposed by Chen et al. (2016), learns the interpretable and disentangled representations in a completely unsupervised manner by maximizing the mutual information between conditional and generated data, which represents a step forward beyond the cGAN. In InfoGAN, an auxiliary network Q is introduced to learn the latent variable y' , as shown in Fig. 9, where Q takes the output from $G(z|y)$ to find the representation code y' , while minimizing the cross entropy between y and y' . In practice, to reduce computation costs, Q and D share most of the weights except the last fully connected layers. The loss function used in InfoGAN is information regularized, which is defined as follows:

$$L_{\text{InfoGAN}} = L(D, G) - \beta I(G, Q) \quad (7)$$

where $L(D, G)$ is the adversarial loss in the original GAN, $I(G, Q)$ is the mutual information and β is the tunable regularization parameter. Training InfoGAN aims to maximize the mutual information between the feature representation delivered on the input of the discriminator and the generated image. InfoGAN delivered impressive performance on different image datasets (Chen et al., 2016). There are several variants of InfoGAN, such as semi-supervision InfoGAN (Spurr et al., 2017) and causal InfoGAN (Kurutach et al., 2018).

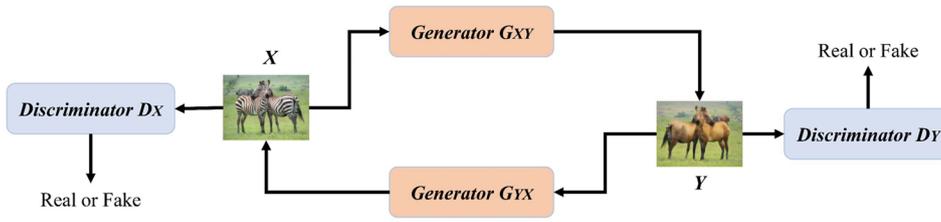


Fig. 8. The framework of cycle generative adversarial network (CycleGAN) that contains four networks, i.e., two discriminators, D_X and D_Y , and two generator networks, G_{XY} and G_{YX} .

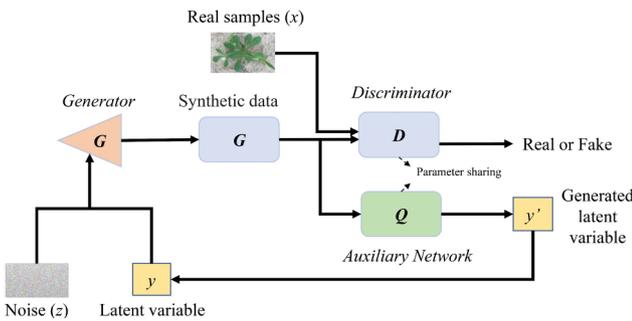


Fig. 9. The framework of information maximizing generative adversarial network (InfoGAN).

4.6. Super resolution generative adversarial network

Image super-resolution aims to recover a high-resolution image upsampled from a low-resolution image, which is a fundamentally challenging task because of the ill-posed nature of the problem (Nasrollahi and Moeslund, 2014). Super resolution generative adversarial network (SR-GAN) (Ledig et al., 2017) represents the first framework to infer photo-realistic natural images for up to $4 \times$ upscaling factors by adopting the GAN framework, where the discriminator is trained to differentiate the synthesized super-resolved images from the original images. The SR-GAN model is built with residual blocks (He et al., 2016) as a generator and the VGG network (Simonyan and Zisserman, 2014) as a discriminator, and it is optimized using a new perceptual loss function that is formulated as a weighted sum of adversarial loss and a content loss (Ledig et al., 2017). The perceptual loss is defined on high-level feature maps of the VGG network that can be invariant to changes in pixel space. SR-GAN showed dramatic improvements in perceptual quality over state-of-the-art image reconstruction methods (Ledig et al., 2017). An enhanced SR-GAN (ESRGAN) (Wang et al., 2018b), significantly improves image sharpness and textural details through three major modifications made to SR-GAN. A new residual-in-residual dense block (RRDB) with BN layers removed was used as the network building block, and the relativistic average discriminator (RaD) (Jolicœur-Martineau, 2018) was used to predict relative realness instead of the absolute value, and the perceptual loss was improved by using features before instead of after activation as practiced in SR-GAN. Later, Real-ESR-GAN (Wang et al., 2021), trained with pure synthetic data, further extends ESR-GAN for restoring real-world degraded images.

4.7. ProGAN

Generating high-resolution/-fidelity images remains a challenging task for GANs, since higher resolution makes it easier for the discriminator to distinguish between generated and real samples (Odena et al., 2017). Progressively Growing GAN (ProGAN) (Karras et al., 2018) employs the idea of progressive neural networks originally proposed in Rusu et al. (2016). It involves progressively growing the generator and discriminator networks in synchrony and training the models starting

from low-resolution (4×4 pixels) incrementally to high-resolution images by adding layers to the networks, as illustrated in Fig. 10. This strategy enables significantly improving training speed (2–6 times speedup) and stability at large image sizes. The incremental expansion on the convolutional layers allows the generator and discriminator models to effectively learn coarse-level details at the beginning and later finer-scale detail as training progresses. Such a progressive training strategy is also employed by other advanced GANs such as StyleGAN (Karras et al., 2019) and BigGAN (Brock et al., 2019), for synthesizing plausible images. In ProGAN, the authors also applied skip connections to add new layers either to the output of the generator or the input of the discriminator through a weighting scheme to control the influence of the newly added layers. During the training, an equalized learning rate was applied to all parameters, and pixelwise feature vector normalization was done to prevent the escalation of parameter magnitudes, and WGAN-GP (gradient penalty) loss (Gulrajani et al., 2017) was used for model optimization. ProGAN was able to produce high-quality images with resolutions up to 1024×1024 pixels and achieve a record inception score (IS) (Salimans et al., 2016) of 8.8 for CIFAR10 (10 classes of 32×32 RGB images) (Krizhevsky and Hinton, 2009).

4.8. Style generative adversarial networks

Despite developments in GAN architectures, the generator remains a black box in image synthesis (Bau et al., 2018). Motivated by image style transfer (Huang and Belongie, 2017), StyleGAN introduced by Karras et al. (2019), is featured with a re-designed generator architecture that enables scale-specific control over image synthesis. StyleGAN adopts the baseline of ProGAN architecture (Karras et al., 2018) and replaces the original noise vector input of the generator with two new references of randomness to produce a synthetic image: standalone mapping channels and noise layers as shown in Fig. 11, where an adaptive instance

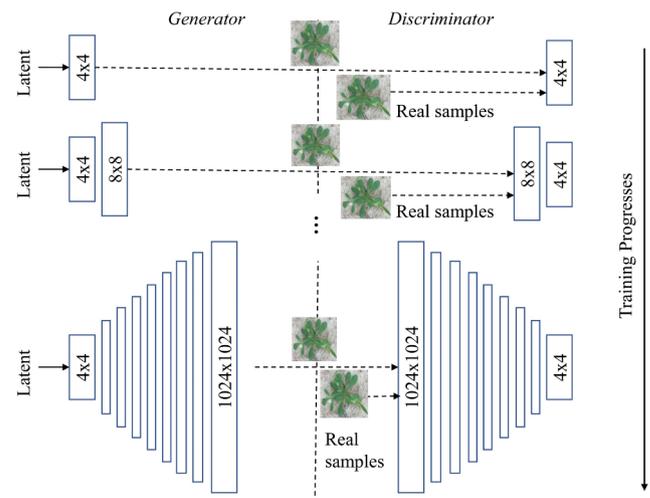


Fig. 10. The architecture of progressively growing GAN (ProGAN). Adapted from Karras et al. (2018)

normalization (AdaIN) layer (Huang and Belongie, 2017) is employed to control the characteristic of the generated images. StyleGAN employs two random latent codes in the mapping function and a split point in the synthesis network. All AdaIN operations prior to the split points use the first style latent vectors, and all AdaIN operations after the split points get the second style latent vectors, which is referred to mixing regularization. This encourages the layers and blocks to localize the style to specific parts of the model and corresponding level of details in the generated image. Stochastic variation is introduced through adding Gaussian noise in the *generator* model and broadcasted to all feature maps, which enables the generation network to synthesize images diverse in fine-grained, per-pixel manner. In Karras et al. (2019), StyleGAN achieved about 20% improvements over the ProGAN in terms of Fréchet inception distance (FID) (Heusel et al., 2017). Built upon the StyleGAN, two improved versions, StyleGAN2 (Karras et al., 2020b) and StyleGAN3 (Karras et al., 2021), have been developed to remove image artifacts (e.g., texture sticking, unwanted information leaked into the image synthesis process) for better quality of generated images.

4.9. Self-Attention generative adversarial network

CNNs are commonly used to build GANs for image generation. Since CNNs generally process information in a local neighborhood, convolutional GANs have difficulty in modeling long-range dependencies beyond local regions. They may fail to capture geometric or structure patterns that occur consistently in some classes when learning multi-class image datasets. Zhang et al. (2019a) proposed the Self-Attention GAN (SA-GAN) for modeling long-range dependencies for image generation tasks. SA-GAN incorporates a self-attention mechanism (Vaswani et al., 2017), as shown in Fig. 12, into the convolutional GAN framework. Adapted from the nonlocal model of Wang et al. (2018a), the self-attention module is deployed in the design of the *generator* and *discriminator* architectures (Zhang et al., 2019a), which is complementary to convolutions for image generation. SA-GAN enables a large receptive field for CNNs to model long-range, multi-level dependencies across image regions, without sacrificing computational efficiency. In SA-GAN, spectral normalization (SN) (Miyato et al., 2018) is applied to the *generator* for stabilized training. Experimented on the ImageNet dataset, SA-GAN achieved state-of-the-art performance on multi-class image synthesis (Zhang et al., 2019a).

4.10. BigGAN

BigGAN (Brock et al., 2019) was designed for scaling up GAN models for generating high-resolution, high-fidelity images through a regularization scheme. Built on the SA-GAN architecture (Zhang et al., 2019a), BigGAN employs a set of regularization techniques for scaling-up image synthesis, such as moving average of generator's weights, class-conditional BatchNorm (Dumoulin et al., 2016), orthogonal

initialization scheme, direct skip connections (skip- z), larger model and batch size, truncation tricks, etc. Particularly, BigGAN adopts a shared embedding approach in the BatchNorm layers to significantly reduce computation costs and improve training speed. The authors successfully trained the model on ImageNet images at 128×128 , 256×256 and 512×512 pixel resolutions (Brock et al., 2019), improving on the state of the art by a large margin. It is demonstrated that GANs benefit from larger models (4 times bigger) and larger batch sizes (8 times larger), since large networks and batch size can significantly increase the capacity of models to handle complex datasets like ImageNet.

4.11. Transformer generative adversarial network

While widely used in modern GAN architectures, CNNs are considered not well suited to process long-range dependencies because of limited local receptive fields, which may lead to loss of feature resolution and fine details as well as the difficulty of model optimization. To address these potential issues, Jiang et al. (2021) proposed a novel transformer-based GAN architecture completely free of convolutions, called TransGAN. Transformer (Han et al., 2022) is an emerging family of DL methods based on self-attention mechanisms, initially applied to natural language processing and subsequently spread in computer vision applications. Compared to CNNs, transformers are advantageous as they enable modeling long dependencies between input sequence elements and support parallel processing of sequence. In TransGAN, the pure transformer-based *generator* and *discriminator* are used in place of CNN as GAN building blocks. For memory efficiency, the *generator* consists of multiple transformer modules, which progressively increases the feature map resolution, as shown in Fig. 13. Furthermore, a multi-scale structure *discriminator* is used to capture semantic contexts and low-level textures simultaneously. Training techniques including data augmentation and modified layer normalization are adopted to stabilize model optimization and generalization. Tested on low-resolution (e.g., 48×48 and 128×128 pixels) benchmark datasets, TransGAN outperformed several state-of-the-art CNN-based GANs (Jiang et al., 2021), such as WGAN-GP (Gulrajani et al., 2017), SN-GAN (Miyato et al., 2018), AutoGAN (Gong et al., 2019) and StyleGAN-V2 (Karras et al., 2020b). However, the performance of TransGANs in synthesizing high-resolution images (e.g., 1024×1024 pixels) remains to be improved or extensively assessed.

4.12. Summary

Table 2 summarizes various GAN architectures developed for image synthesis. These architectures do not represent the full landscape of GANs in literature, but they, coming with open-sourced packages, have been employed or adapted for expanding agricultural image data, as reviewed in Section 5. In addition to the tabular summary, readers are referred to generic GAN literature (Creswell et al., 2018; Liu et al., 2021; Huang et al., 2018) and “The GAN Zoo” in the GitHub (<https://github.com/hindupuravinash/the-gan-zoo>) for a more comprehensive list of GAN variants as well as a GAN toolbox in the PyTorch environment at https://github.com/facebookresearch/pytorch_GAN_zoo for algorithm implementation.

5. Applications of GANs in agriculture

Presented in this section is a summary review of application studies of GAN in the agricultural field and food domain. These applications are loosely organized into three areas, including precision agriculture (Table 3), plant phenotyping (Table 4), and postharvest quality assessment of agricultural products (Table 5).

5.1. Precision agriculture

Precision agriculture, also known as precision or smart farming, aims to enhance crop and animal production using more precise (e.g., site-

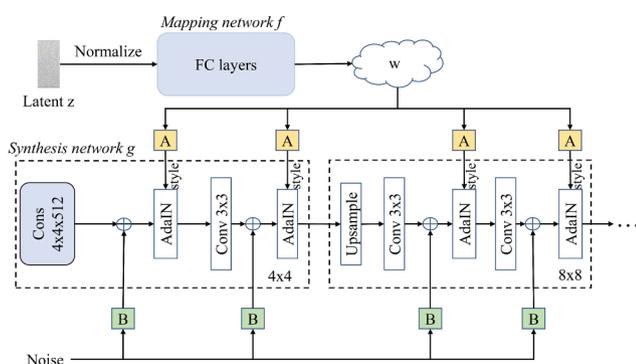


Fig. 11. The architecture of style generative adversarial network. Adapted from Karras et al. (2019)

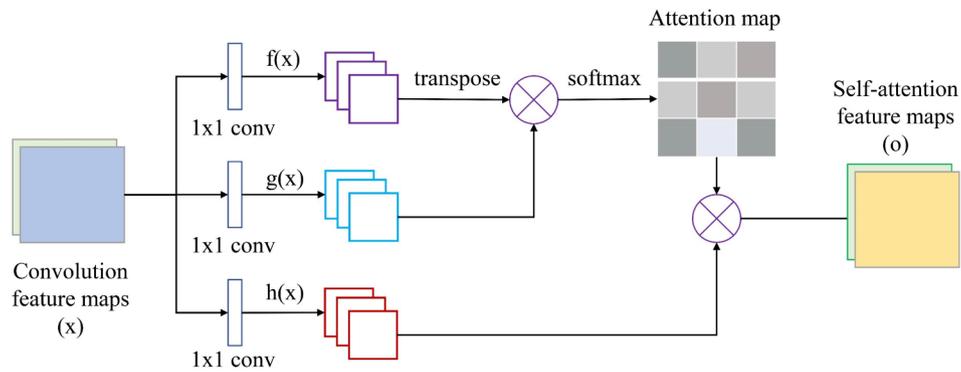


Fig. 12. The architecture of the self-attention network used in self-attention generative adversarial network. Adapted from Zhang et al. (2019b)

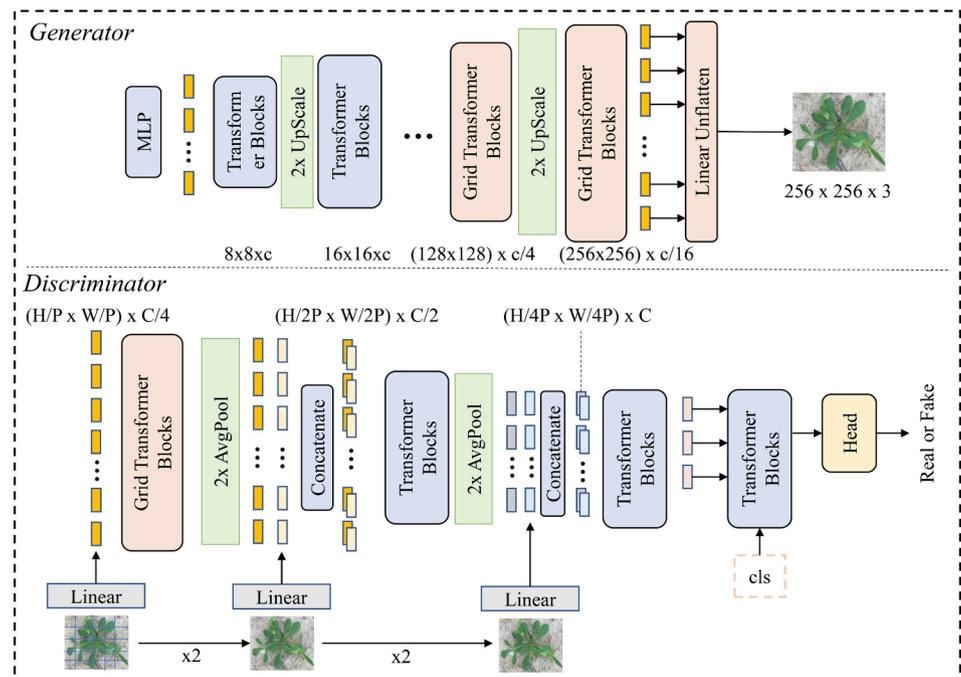


Fig. 13. The framework of a transformer generative adversarial network. Adapted from Wang & Xiao (2021)

specific) and resource efficient approaches (Monteiro et al., 2021). It exploits an assortment of monitoring and invention technologies (e.g., imaging, autonomy, AI/robotics) to support data driven agricultural tasks, such as crop health detection, weed recognition and control, and animal monitoring and tracking. Over the past few years, GANs have been increasingly applied to augment image data in precision agriculture and enhance machine learning models for various applications as reviewed below.

5.1.1. Plant health

Biotic and abiotic stressors (Dhaka et al., 2021), such as microorganisms (e.g., viruses, bacteria, fungi), insects, and environmental factors, negatively affect plant growth and health conditions, leading to the development of plant diseases or disorders and eventually low yield and quality of plant products (Zhang et al., 2021b). Imaging technologies (e.g., RGB, NIR, fluorescence, thermal) offer a non-invasive and objective means for characterization and diagnosis of plant health conditions (Thomas et al., 2018; Mahlein et al., 2018). One of the grand challenges in imaging-based plant health detection is how to cope with limited image data with a small number of expert annotations, especially for

emerging plant diseases for which it can be very difficult to accumulate large amounts of image data that are ground-truthed with specialized knowledge. This in turn hampers the implementation of supervised machine learning methods. The developments of GANs for image synthesis, given successful applications in medical diagnosis (Yi et al., 2019), has inspired significant efforts made to utilize GANs to augment image data for enhancing the detection of plant diseases and other health conditions.

CGAN (Section 4.1) / DCGAN (Section 4.2) have received a great deal of attention in image augmentation for improved plant disease recognition. Hu et al (2019) reported on using DCGAN with conditional label constraint (named C-DCGAN) for identifying tea leaf diseases (red scab, red leaf spot and leaf blight). Original color images, subjected to disease segmentation by support vector machine (SVM), were augmented from 20 to 4980 images per disease and fed into classifiers (VGG16, SVM, decision tree and random forest). With the GAN-augmented data, VGG16 yielded the best average accuracy of 90%, which was 28% higher than that using rotation and translation-based augmentation and even higher than the accuracy by VGG16 without data augmentation due to severe overfitting. Abbas et al. (2021) used

Table 2
Summary of GAN variants with open-source software packages.

GAN architecture	URL	Reference
AutoGAN	https://github.com/VITA-Group/AutoGAN	Gong et al. (2019)
BigGAN	https://github.com/ajbrock/BigGAN-PyTorch	Brock et al. (2019)
CGAN/ CDCGAN	https://github.com/znxlwm/tensorflow-MNIST-cGAN-cDCGAN	Mirza & Osindero (2014)
CoGAN	https://github.com/mingyuliutw/cogan	Liu & Tuzel (2016)
CycleGAN	https://github.com/junyanz/pytorch-CycleGAN-and-pix2pix	Zhu et al. (2017)
DCGAN	https://github.com/eriklindernoren/PyTorch-GAN/blob/master/implementations/dcgan/dcgan.py	Radford et al. (2015)
ESR-GAN	https://github.com/xinntao/ESRGAN	Wang et al. (2018a,b)
InfoGAN	https://github.com/openai/InfoGAN	Chen et al. (2016)
Vanilla GAN	https://github.com/goodfeli/adversarial	Goodfellow et al. (2014)
Pix2Pix	https://github.com/phillipi/pix2pix	Isola et al. (2017)
ProGAN	https://github.com/tkarras/progressive_growing_of_gans	Karras et al. (2018)
SA-GAN	https://github.com/brain-research/self-attention-gan	Zhang et al. (2019b)
StyleGAN	https://github.com/NVLabs/stylegan	Karras et al. (2019)
StyleGAN2	https://github.com/NVLabs/stylegan2	Karras et al. (2020b)
StyleGAN3	https://github.com/NVLabs/stylegan3	Karras et al. (2021)
StyleSwin	https://github.com/microsoft/StyleSwin	Zhang et al. (2021a)
SR-GAN	https://github.com/tensorlayer/srgan	Ledig et al. (2017)
TransGAN	https://github.com/VITA-Group/TransGAN	Jiang et al. (2021)
WGAN	https://github.com/eriklindernoren/Keras-GAN	Arjovsky et al. (2017)
WGAN-GP	https://github.com/caogang/wgan-gp	Gulrajani et al. (2017)

CGAN (Section 4.1) to augment tomato leaf images of ten classes (Fig. 14), curated from the Plant Village dataset (Hughes et al., 2015). With 4000 synthetic images per class added to training data, classification accuracies of over 97% were achieved by DenseNet121, which were about 1–4% higher than the accuracies on the original data without image augmentation. However, the paper did not investigate the effect of different number of synthetic images on the developed classification model.

Douarre et al. (2019) presented an early study using DCGAN to generate images for segmenting apple scab (a disease manifested as spot lesions on the leaves and fruits). The infrared (IR) images acquired for plant canopy were titled into sub-images of 64 × 64 pixels and fed into SegNet (Badrinarayanan et al., 2017) for pixelwise segmentation (“scab” vs “not scab”). Compared to a baseline model without involving data augmentation, DCGAN with the WGAN training process (Arjovsky and Bottou, 2017) yielded a 2% improvement in segmentation accuracy. Larger improvements were, however, achieved using standard data augmentation and model-based simulation approach (Douarre et al., 2019), because DCGAN was found difficult to converge to produce convincing images. Zeng et al. (2020) curated a citrus leaf dataset comprising 5406 color images with visual symptoms of Huanglongbing (a citrus greening disease) from the Plant Village dataset (Hughes et al., 2015). DCGAN was adapted to expand the training data (80% of the total), synthesizing additional 8650 images (doubling the number of training samples). An improvement of up to 20% classification accuracy was achieved using the GAN-augmented data. Zhang et al. (2019b) also employed DCGAN to expand images of citrus canker disease and

Table 3
Various GAN methods applied for plant health detection. Acc and F1 denote overall classification accuracy and F1-score, respectively.

GAN Architecture(s)	Plant Materials	Reference	Performance Metrics
CGAN/CDCGAN	Tea leaves	Hu et al. (2019)	Acc gains 28%
	Tomato leaves	Abbas et al. (2021)	Acc = 97%
DCGAN	Apple scab	Douarre et al. (2019)	Acc gains 2%
	Tomato leaves	Wu et al. (2020); Hu et al. (2021)	Acc = 79%–94%
	Citrus canker	Zhang et al. (2019b)	Acc = 90.1%
	Citrus leaves	Sun et al. (2020)	Acc = 97.8%, gains 2.3%
		Zeng et al. (2020)	Acc gains 20%
	Tea leaves	Yuwana et al. (2020)	Acc = 88.84%, gains 2.5%
	Blueberry leaves	Kim et al. (2021)	Acc = 97.3–98.8%
CycleGAN	Tomato leaves	Nazki et al. (2019)	–
	Apple leaves	Tian et al. (2019)	F1 gains 5%
	Cucumber leaves	Cap et al. (2020)	Acc gains 7.4%
	Grape leaves	Zeng et al. (2021)	Acc gain 7%
Info-GAN	Tomato leaves	Gomaa & El-Latif (2021)	Acc = 97.9%
SR-GAN	Wheat stripe rust	Maqsood et al. (2021)	Acc = 83%, gains 8%
ESR-GAN	Tomato leaves	Wen et al. (2020)	Acc = 90.8%, 18.0% gains
StyleGAN	Multiple species	Arsenovic et al. (2019)	Acc = 90.9%, 0.9% gains
Others	AC-GAN	Lu et al. (2019)	F1 = 0.95, gains 0.3
	Activation reconstruction GAN	Nazki et al. (2020)	Acc gains 5.2%
	Channel decreasing GAN	Liu et al. (2020)	Acc = 98.7%
	PSR-GAN	Dai et al. (2020a)	Acc = 89.1–94.2%
	DATF-GAN	Dai et al. (2020b)	Acc = 88.1–92.7%
	WGAN-GP	Bi & Hu (2020)	Acc gains 24.4%
	DoubleGAN	Zhao et al. (2021b)	Acc = 99.8%, 99.5%
	RAHC_GAN	Deng et al. (2021)	Acc = 98.0–99.8%
	Re-enforcement GAN	Nerkar & Talbar (2021)	Acc = 99.0%
	Style consistent image translation	Xu et al. (2022)	Acc = 73.1–96.5%

obtained a classification accuracy of 90.1%. A further study was conducted on generating lesion images with a specific shape by feeding binary images into the GAN generator (Sun et al., 2020), which improved the recognition of leaf lesion from 95.5% on the original data to 97.8% for synthetic images.

Yuwana et al. (2020) investigated the vanilla GAN (with multilayer perceptron as generator and discriminator models) and DCGAN to synthesize tea leaf images of four classes (healthy and three diseased) for disease identification. With 1000 images synthesized per class, the vanilla GAN and DCGAN yielded classification accuracies of 88.84% and 88.86%, respectively, representing about 2.5% improvements over the

Table 4

Applications of GANs in precision/production agriculture (see also Table 3). Acc, IoU, F1 and mAP denote overall classification accuracy, intersection over union, F1 score, and mean average precision (for object detection), respectively.

Application	GAN architecture	Materials examined	Reference	Performance Metrics
Weed Control	CGAN/CDCGAN	Sugar beet, sunflower	Fawakherji et al. (2021)	Acc = 94.6%, IoU = 96.6%
	DCGAN	Tomato, black nightshade	Espejo-Garcia et al. (2021)	F1 = 99.1%, gains 0.85%
	Semi-Supervised GAN	Sugar beet, weeds	Kerdegari et al. (2019)	F1 = 85%
	DCGAN	Pea, strawberry, weeds	Khan et al. (2021)	Acc = 90.0%
Fruit Detection	CycleGAN	Sweet/bell pepper	Barth et al. (2020)	IoU gains 8%
	BEGAN	Pine cones	Luo et al. (2020)	AP = 93.4% – 95.3%
	CycleGAN	Orchard	Bellocchio et al. (2020)	–
	CycleGAN	Vineyard	Fei et al. (2021)	mAP@IoU0.5 = 37.20%
	CGAN	Kiwi fruits	Olatunji et al. (2020)	Error = 5%-7%
Aquaculture	CGAN/CDCGAN	Grapevine berries	Kierdorf et al. (2021)	Correlation = 0.75–0.98
	DCGAN	Fish species	Zhao et al. (2018)	Acc = 80.5%–83.1%
	CycleGAN	Shrimp eggs	Zhang et al. (2021c)	Acc = 99.2%
Animal Farming	SAGAN	Goats	Li & Tang (2020)	Acc = 85.1%
	CGAN	Cattle muzzles	Singh et al. (2021)	Acc = 98.9%
	CTGAN	Poultry chickens	Ahmed et al. (2021)	Acc = 97.0%

baseline models without image augmentation. However, generating more training samples led to mixed performance, probably due to the inclusion of more low-quality (e.g., noisy, texture detail loss) images generated by the GANs. Wu et al. (2020) applied DCGAN to generate images of tomato leaves of 5 classes (healthy and four diseased). A set of 1500 images extracted from the Plant Village dataset (Hughes et al., 2015) were augmented to 5300 images to train CNN classification models, resulting in accuracies of 79%–94%. However, the authors did not systematically compare models trained with and without using DCGAN. In classifying images of seven tomato leaf diseases, Hu et al. (2021) used DCGAN to augment image data for different CNN classifiers (VGG19, AlexNet and ResNet50) and achieved accuracy improvements of 1% to 3%, compared to models without image augmentation. In using

Table 5

Application of GANs in plant phenotyping and post-harvest quality assessment. R2, Acc, F1, SSIM and mAP denote coefficient of determination (for regression models), overall classification accuracy, F1 score, structural similarity index, and mean average precision, respectively.

Application	GAN architecture	Materials examined	Reference	Performance Metrics
Plant Phenotyping	DCGAN	Arabidopsis plants	Giuffrida et al. (2017)	R ² = 0.95
	CGAN (Pix2Pix)	Arabidopsis plants	Zhu et al. (2018)	Error reduction 16.7%
	WacGAN	Multiple species plant seedlings	Madsen et al. (2019a)	Acc = 58.9%
	WacGAN-Info	Multiple species plant seedlings	Madsen et al. (2019b)	Acc = 64.3%
	CDCGAN	Orchid vigor rating	Zhu et al., (2020)	F1 gains 23%
	CGAN (Pix2Pix)	Arabidopsis plants and field-grown cauliflower	Drees et al. (2021)	R ² = 0.66–0.95
	TasselGAN (DCGAN)	Maize tassels	Shete et al. (2020)	SSIM = 0.84–0.89
	CycleGAN	Wheat heads	Hartley and French (2021)	–
Postharvest assessment	Customized GAN	Coffee beans	Chou et al. (2019)	Acc = 80%
	DCGAN	Jujubes	Guo et al. (2021)	Acc = 98%
	CGAN	Lemons	Bird et al. (2022)	Acc = 88.8%, gains 5.0%
	CoGAN	Potatoes	Marino et al. (2020)	F1 = 0.84
	TransGAN	Lychees	Wang & Xiao (2021)	mAP = 91.8%
	Vanilla GAN	Wheat kernels	Yang et al. (2021)	F1 = 96.2%

DCGAN to generate images for early identification of tomato plants infected tomato mosaic virus, Gomaa & El-Latif (2021) obtained the recognition accuracy of 97.9% using the augmented data, about one percent better than the accuracy without augmentation.

CycleGAN (Section 4.4) has also been evaluated to synthesize and augment image data in plant health detection. Nazki et al. (2019) applied CycleGAN with U-Net (Ronneberger et al., 2015) as a generator for image synthesis to rebalance a 9-class tomato disease dataset. In terms of perceptual image quality metrics such as FID, CycleGAN with U-Net performed better than the classic CycleGAN in capturing low-level details and realistic texture. However, this study did not examine the performance of synthesized images in plant disease recognition tasks. In applying CycleGAN to a small 4-class grape disease dataset, Zeng et al. (2021) obtained accuracy improvements of over 7% with ResNet18, compared to the model on the original data. Tian et al. (2019) used CycleGAN to complement traditional data augmentation methods for enhanced detection of anthracnose lesions on apple fruits. The dataset, collected from both orchards and online, consisted of 140 diseased and 500 healthy apple images, and CycleGAN was applied to transform healthy apple images into diseased fruit images. With the inclusion of CycleGAN-synthesized images, the disease detection model, based on YOLOv3 (Redmon and Farhadi, 2018), gained about 5% improvements in terms of F1 score and IoU (intersection over union). Despite these applications, CycleGAN has limitations in generating high-quality images because of no explicit attention mechanisms for transforming specific objects in images, and consequently the generated images may contribute little to model robustness in plant disease diagnosis (Cap et al., 2020). To enrich the versatility of image generation, Cap et al. (2020) introduced a leaf segmentation module (composed of a weakly supervised segmentation network) to CycleGAN, leading to a model called LeafGAN, to transform regions of interest in plant disease images. Tested on 5-class cucumber leaf disease data, LeafGAN yielded diagnostic performance improvements by 7.4% as opposed to 0.7% by the vanilla CycleGAN.

High-resolution images are crucial for effective recognition and detection of plant diseases. SR-GAN (Section 4.6) enables synthesizing super-/high-resolution images with photo-realistic details from down-sampled, low-resolution images. Maqsood et al. (2021) examined SR-GAN combined with a CNN model for classifying wheat yellow rust disease images. The original data, consisting of 1922 color images of three classes (healthy, resistant, and susceptible) were first preprocessed to prepare down-sampled, low-resolution images, and then fed into SR-GAN with an up-sampling factor of 4 × to produce photorealistic images. The synthetic images yielded an overall classification accuracy of 83%, which outperforms the down-sampled, low-resolution images that only gave 75% overall accuracy. However, this study did not apply SR-GAN to the original images for image synthesis. Built upon SR-GAN, an enhanced version, i.e., ESRGAN (Wang et al., 2018b), can produce

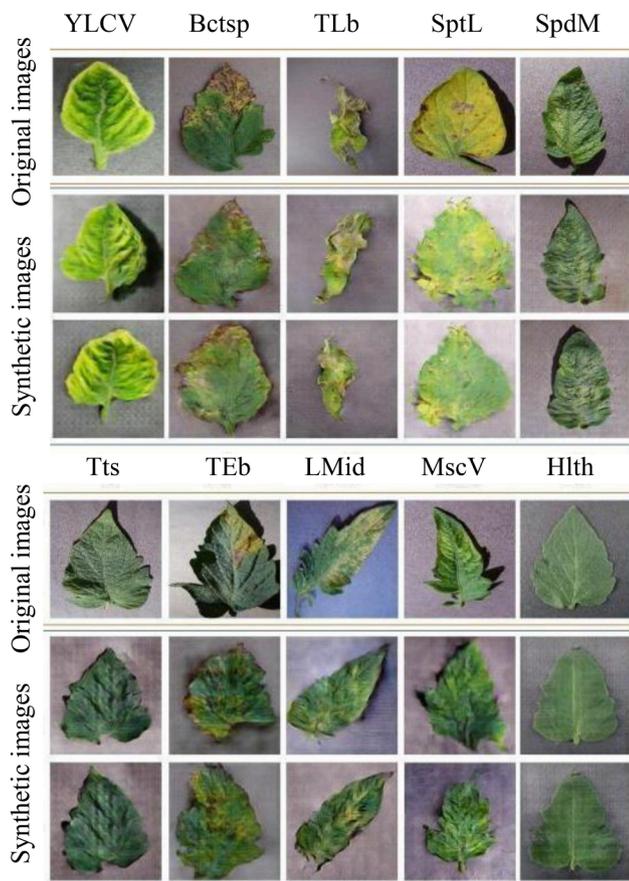


Fig. 14. Original tomato leaf images and the generated images by CGAN. Reproduced with permission from Abbas et al. (2021)

better visual quality with more realistic textures. Wen et al. (2020) reported on using ESR-GAN for image synthesis for a tomato disease dataset extracted from the Plant Village dataset (Hughes et al., 2015). Likewise, ESR-GAN was applied to low-resolution target images with a scaling factor of $\times 4$ and fine-tuned for the domain task. The generated images, trained on VGG16, yielded an accuracy of 90.78% in the 10-class classification, substantially better than 72.74% by using low-resolution images. It is noted that the accuracy obtained by the synthesized SR images is still about 4% lower than that by the original high-resolution images.

StyleGAN (Section 4.8) that combines ProGAN and neural style transfer has also been used to generate high-resolution images for enhanced plant disease recognition. In Arsenovic et al. (2019), StyleGAN along with traditional data augmentation methods (geometric transformations, such as rotations or pixel-wise changes) was employed to augment the Plant Disease dataset (Brahimi et al., 2017), containing 79,265 images of 12 different species and 42 classes (including both healthy and diseased). To compare the approaches while considering the impact of environment conditions, the authors trained five CNN plant disease classification models on the PlantVillage dataset (Hughes et al., 2015) and tested the models on the Plant Disease dataset (Brahimi et al., 2017) with both standard augmentation and the StyleGAN. The models trained on the StyleGAN-augmented dataset (top-1 accuracy of 0.9088 for ResNet-152) outperformed the models trained on the dataset with standard augmentation (top-1 accuracy of 0.8995 for ResNet-152). The authors also reported on leave disease detection, but performance comparisons regarding data augmentation were not detailed. Further research can be done on the application of new-generation StyleGAN variants, such as StyleGAN2 (Karras et al., 2020b) and StyleGAN3 (Karras et al., 2021), for better quality of generated images.

Apart from the aforementioned GAN methods, some other GAN variants have also been used for plant health detection. Lu et al. (2019) reported on using AC-GAN (auxiliary classifier GAN) (Odena et al., 2017) to generate synthetic insect pest images for classifying 5-class insect images and obtained an F1-score of 0.95 (0.3 higher than that based on traditional data augmentation). Nazki et al. (2020) proposed AR-GAN (activation reconstruction GAN) by introducing a module to calculate activation reconstruction loss in CycleGAN (Zhu et al., 2017) to rebalance an imbalanced 9-class tomato disease dataset. The GAN-synthesized data yielded an improvement of 5.2% in classification accuracy as compared to an 0.8% increase with traditional augmentation. The AR-GAN approach was later adopted by Zhang et al. (2021b) for improving the identification of cucumber leaf diseases. Liu et al. (2020) presented a GAN model with a channel decreasing generator to synthesize 4-class grape leaf images, reporting 98.7% classification accuracy, which is about 3% and 2% better than the models without and with only basic image augmentation, respectively. Recently, Xu et al. (2022) adapted style consistent image translation GAN (SCITGAN) (Wang et al., 2019) for tomato disease recognition. A five-class tomato leaf dataset was collected for GAN training to generate images (Fig. 15) and instances for each disease class. Among different object detection and instance segmentation models, PointRend (Kirillov et al., 2020) trained with the SCITGAN-augmented data achieved the best segmentation accuracy with mAP (mean average precision) of 68.3%, outperforming the models with CycleGAN (67.6%) and without augmentation (56.1%). More application studies of GANs on plant health detection are given in Table 3.

5.1.2. Weed control

Weeds compete with crops for resources (i.e., water and nutrients) and provide hosts for pests and diseases, posing a vital threat to crop production (Chen et al., 2022). Machine-vision-based weed control offers a promising means for managing weeds effectively, especially mitigating herbicide resistance, by identifying and localizing weed plants followed by site-specific, individualized treatments (e.g., spot spraying, killing weeds by high-flame laser). Currently it remains a critical challenge to develop machine vision systems capable of accurate weed recognition and robust to variable field light conditions (Chen et al., 2022; Westwood et al., 2018).

Fawakherji et al. (2021) applied CGAN (Section 4.1) to generate four-channel multispectral (RGB + NIR) images for crop/weed segmentation. Instead of pursuing a full-scene image generation, the authors conditioned GAN models on the shape/mask of plant regions to produce semi-artificial images by replacing only targeted plant objects in images with synthetic instances, leaving soil backgrounds unchanged. A spatially adaptive denormalization GAN architecture that synthesizes photorealistic images given semantic layout input (Park et al., 2019) was adapted to perform semantic image synthesis. Evaluation on three public datasets, i.e., two sugar beet datasets (Chebrolu et al., 2017) and the sunflower dataset created by the authors, showed that the GAN-augmented data statistically improved pixel-wise segmentation accuracies compared to the original data, with improvements in mean IoU (mIoU) of 0.03–0.10 depending on datasets and segmentation networks. The ensemble of RGB and NIR images yielded consistently better performance than using RGB conventionally. Espejo-Garcia et al. (2021) used DCGAN (Section 4.2) to generate synthetic tomato and black nightshade images for weed identification tasks on the early crop dataset (<https://github.com/AUAGroup/early-crop-weed>). The FID metric (Heusel et al., 2017) was employed to determine the best DCGAN configuration for image generation. Synthetic tomato and black nightshade images generated were added to the original images and transfer learning was adopted to train several CNN classifiers (e.g., Xception, DenseNet and Inception). The Xception network yielded an optimal F1-score of 99.07% using GAN-augmented training dataset compared to 98.22% of training with traditional augmentation methods.

Unmanned aerial vehicle (UAV) or drone-based imagery has been

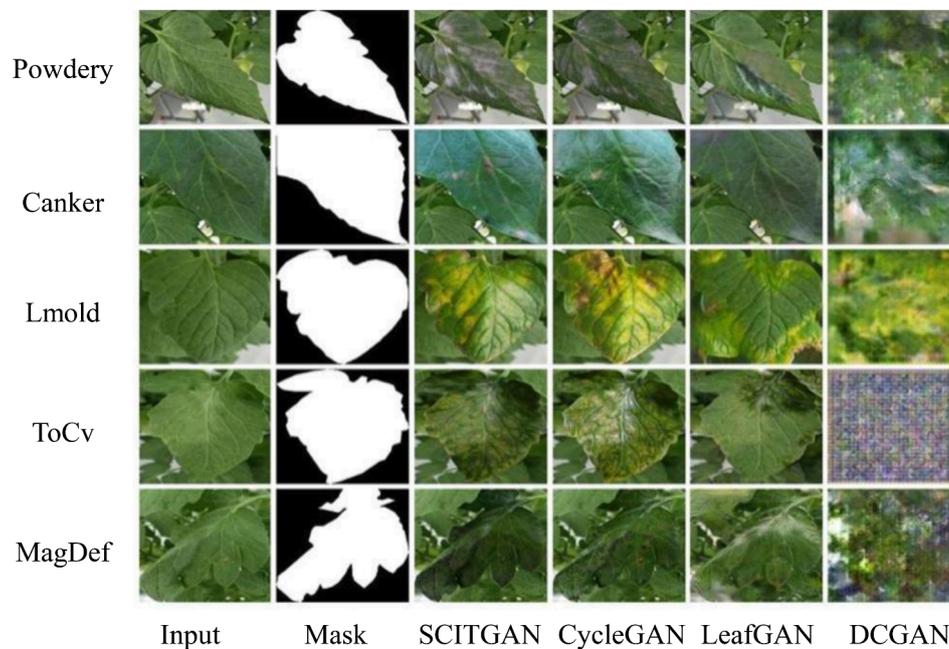


Fig. 15. Sample examples generated by different GAN models (CycleGAN, LeafGAN, DCGAN) against the Style consistent image translation GAN (Xu et al., 2022).

widely used in precision farming to detect plant health conditions including weed infestations. Remote sensing by UAVs coupled with DL algorithms enables field-scale weed recognition and monitoring to support site-specific management. To address data insufficiency in weed-crop classification tasks, Kerdegari et al. (2019) applied a semi-supervised GAN (SS-GAN) model adapted from DCGAN for pixel-wise classification of multispectral images acquired by a UAV. Unlike typical GAN where the *discriminator* is a binary classifier for discriminating real and fake images, the SS-GAN used a multi-class classifier discriminator taking the input of synthetic and label samples as well as unlabeled. Evaluated on the weedNet (Sa et al., 2017), a multispectral (RGB + NIR) dataset acquired by a UAV from sugar beet fields, the GAN model achieved F1 scores of about 0.85 by using two-channel (Red + NIR) images with 50% labeled data (Kerdegari et al., 2019). Khan et al. (2021) employed SS-GAN for classifying crops and weeds in the RGB images acquired by UAVs from pea and strawberry fields at early crop growth stages. The SS-GAN achieved an overall classification accuracy of about 90% when only 20% labeled samples were used for training, which compared favorably with the results obtained by conventional supervised classifiers (e.g., K-nearest neighbor, SVM and CNNs).

5.1.3. Fruit detection

In-orchard fruit detection is challenging due to unstructured environments and variable field light. Over the past few years, DL-based object detectors and segmentation networks have been widely researched for fruit detection towards robotic harvesting, fruit counting and yield estimation (Koirala et al., 2019; Maheswari et al., 2021). To improve fruit recognition performance, GANs (Table 4) have been utilized to generate realistic images for training DL models while reducing the requirements of manual efforts for data collection and labeling.

Barth et al. (2020) used CycleGAN to perform domain adaptive semantic segmentation to support robotic perception. The Capsicum annum dataset (Barth et al., 2018), consisting of 50 images of pepper plants collected in a greenhouse and 10500 synthetic images rendered in Blender (open-source 3D modeling software), was used to train CycleGAN for domain translation (synthetic domain to empirical domain). The mean color distribution correlation with the empirical data was improved from 0.62 (on empirical data) to 0.90 (after translation). Using translated synthetic images for part segmentation gave an 8% IoU improvement by fine-tuning with empirical images compared to only

using synthetic images, and a 55% improvement compared to training without empirical fine-tuning. Luo et al. (2020) reported on the detection of pinecones using boundary equilibrium GAN (BEGAN) (Berthelot et al., 2017) and YOLOv3 (Redmon and Farhadi, 2018). The BEGAN uses an equilibrium enforcing method to control the tradeoff between image diversity and visual quality. The authors used the GAN-generated pinecones instances to replace pinecones in the original images, which resulted in the detection average precision (AP) increasing to 95.3% from 93.4% for the model without image augmentation.

Bellocchio et al. (2020) used CycleGAN to perform image domain translation coupled with weakly supervised learning for fruit counting where only fruit presence-absence labels were given. Experiments on four orchard datasets showed CycleGAN could adapt image conditions (e.g., illumination, blur, saturation) to the target domain and transform fruit colors and textures from a source orchard to a different target orchard. Fei et al. (2021) used CycleGAN with semantic constraints for domain adaptation of grape images while preserving spatial semantics such as fruit position and size (Fig. 16). The GAN model was trained to adapt source domain vineyard images that were 3D rendered by Helios (Bailey, 2019), to day and night images (target domains) collected from the orchard. Fruit detection models with YOLOv3 yielded 37.2 mAP@IoU0.5 by fine-tuning the detection model with the GAN synthesized images for day domain, compared to 37.0 mAP@IoU0.5 without the image synthesis, and 26.7 AP@IoU0.5 compared to 22.8 AP@IoU0.5 for night domain. Although realistic images were generated by the GAN model, further validations are still needed on using physically collected real-world data as source domain images.

Fruit occlusions in canopy pose great challenges to accurate fruit detection/localization and crop yield estimation. Olatunji et al. (2020) applied CGAN to reconstruct complete kiwifruit surfaces by translating occluded fruit into non-occluded images to address partial occlusion problems in fruit detection. By reconstructing missing surface information of occluded fruit, the fruit shape, surface area and weight could be predicted from the reconstructed images. The CGAN was trained on 3000 image pairs computationally created in Blender, for generating the complete fruit surface topography. Experimental results on the synthetic validation dataset show that 75% of the volume prediction had an error rate less than 5% and only 7% of volume predictions had error rate larger than 10%. The CGAN model was able to generate high quality images and yielded a predicted fruit weight error of 5.58% validated on

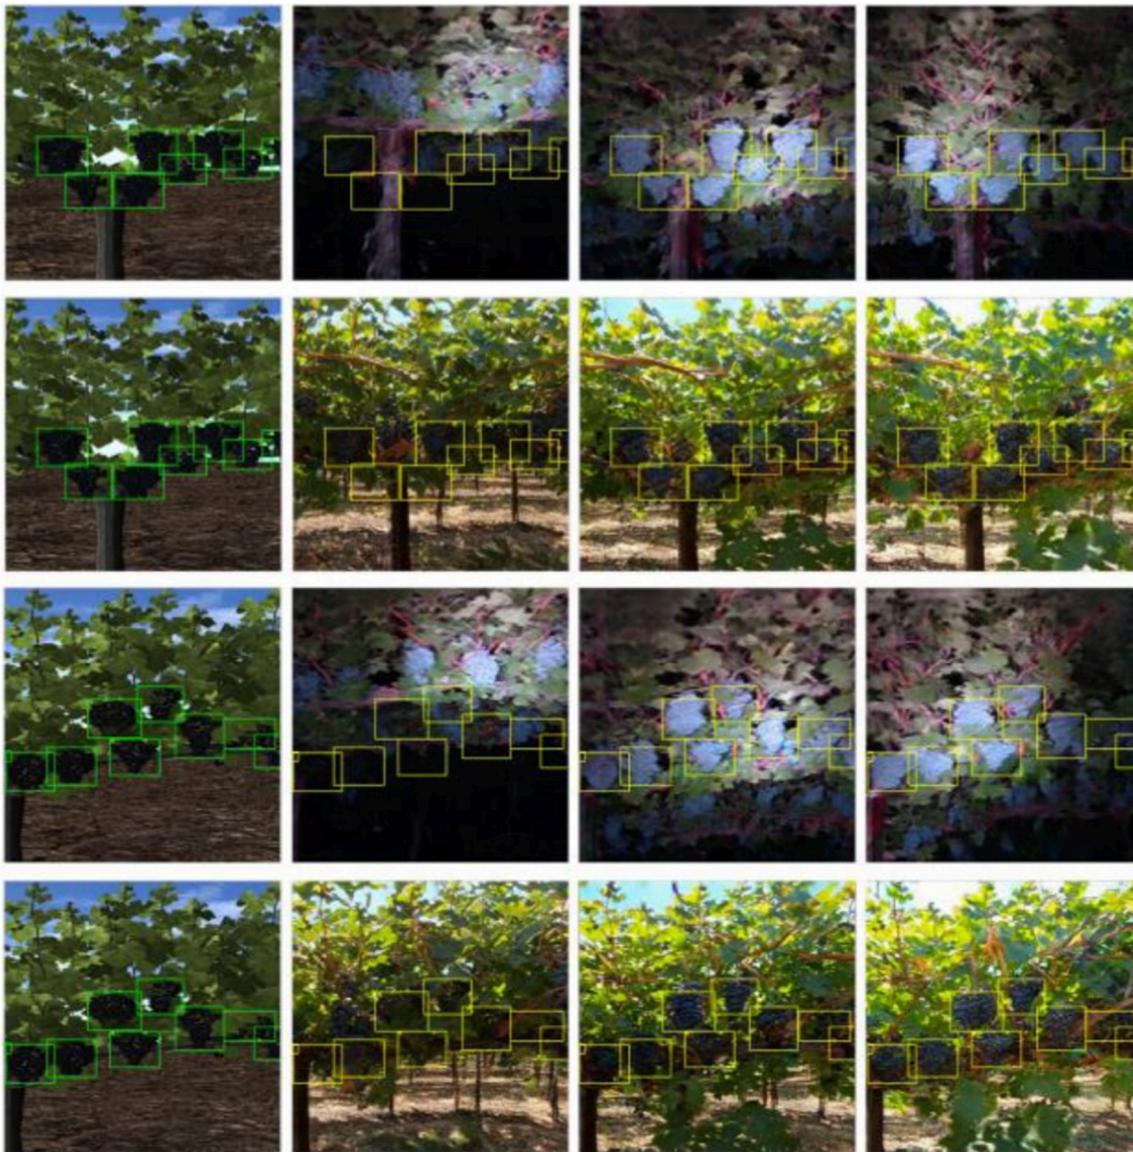


Fig. 16. Images generated by semantically constrained CycleGAN. Column#1 shows 3D rendered source domain images with ground truth labels, Column# 2–4 illustrates images with projected labels highlighted in yellow box. Reproduced with permission from Fei et al. (2021).

a dataset collected by a 3D scanner posed at a single angle to mimic occlusion. The authors however did not discuss the impact on model performance with different numbers of generated training images. Similar to Olatunji et al. (2020), Kierdorf et al. (2021) applied Pix2Pix (Isola et al., 2017) to translate occluded berry images due to leaves to non-occluded fruit images for estimating the number of berries, enabling better yield estimation. More research is needed to use GANs to generate synthetic images for many other specialty crop products for improved fruit detection and yield estimation.

5.1.4. Aquaculture

Application of imaging technology in precision aquaculture faces a number of special challenges (Li et al., 2021a; Føre et al., 2018) in detecting and monitoring aquatic species due to adverse poor underwater conditions (e.g., poor illumination and object visibility in turbid water, cluttered background), which make acquiring high-fidelity/-contrast images difficult. The insufficiency of aquaculture images available would add to the complexity of underwater species recognition tasks.

GANs have been applied for image synthesis for enhancing visual

recognition in aquaculture (Table 4). Zhao et al. (2018) used a semi-supervised learning model based on the modified DCGAN (with a multi-class classifier in the discriminator) for live fish recognition, on two open-sourced fish datasets, i.e., the Fish4-knowledge (Boom et al., 2012) project (27370 fish images in 23 clusters) and the Croatian fish dataset (Jager et al., 2015) (794 images of 12 fish species). Using small fractions (5–15%) of labeled training samples, the semi-supervised GAN model achieved fish identification accuracies of 80.52–83.07%, exceeding the accuracies obtained by CNN-based models (Marburg and Bigham, 2016; Qin et al., 2016) without data augmentation as well as those based on other GANs (e.g., 48.39–59.44% for the standard DCGAN). Improvements were also observed with higher fractions (50–75%) of the labeled Croatian fish data used for fish identification. Zhang et al. (2021c) used CycleGAN (Zhu et al., 2017) to generate images for automatic shrimp egg counting which otherwise would be labor-intensive in the hatcheries. A shrimp egg dataset was collected and manually labeled, containing 450 color images with about 272,000 egg annotations. Randomly generated binary density maps combined with real images without annotations were used to train CycleGAN to generate photo-realistic images with precise egg annotations (Fig. 17). A

shrimp egg counting network combined with the fully convolutional regression network was proposed to count eggs through regressing the input image into its density map. Pretrained on the synthetic images and finetuned on the collected data, the network achieved the average counting accuracy of up to 99.2%, compared to the accuracy of 96.6% by training the synthetic data alone and 98.9% only on the collected dataset.

5.1.5. Animal farming

Animal farming here is restricted to livestock and poultry farming but excludes aquaculture described above. Computer vision systems assisted by DL algorithms have been investigated in animal farming for phenotyping and behavioral identification and monitoring (Wurtz et al., 2019; Oliveira et al., 2021), at group or individual levels, to enhance production efficiency as well as animal welfare, with applications expanded recently (Junior et al., 2021; Li et al., 2021b). However, the exploitation of DL can be hindered by the drudgery of collecting and labeling large amounts of animal image data. GANs are thus attractive because of their ability to generate images and labels automatically to empower DL models.

Li & Tang (2020) evaluated an improved SA-GAN (Section 4.9) that used a normalized self-attention mechanism and replaced one-hot labels with multi-class labels (colors, backgrounds, and behaviors) to generate high quality dairy goat images (Fig. 18). Experiments on a collection of goat images confirmed the effectiveness of the GAN model for enhancing image quality compared to image generation using one-hot labels. Singh et al. (2021) applied CGAN to generate enhanced images for cattle identification based on muzzle pattern images. Low-quality images that were obtained by artificially degrading the original images, alongside binarized muzzle pattern images (the target enhanced images), were

used to train the GAN model to learn image enhancement mapping. The synthetic images coupled with feature extraction and matching achieved an identification accuracy of 98.86% for four breeds of cattle. This study, however, did not report on the cattle identification accuracy using the original images. Wearable sensors have been widely researched for animal tracking and monitoring the behaviors of individuals (Siegford et al., 2016). In Ahmed et al. (2021), the time-series data collected by activity sensors placed on chicks (Abdoli et al., 2018) was examined for classifying healthy and sick chickens. A GAN method designed for modeling tabular data, i.e., conditional tabular GAN (Xu et al., 2019), was used to expand the original sensor data for the classification, resulting in the accuracy of 97% by TabNet (Siegford et al., 2016). Currently, the applications of GANs are still scant for synthesizing image data in poultry farming.

5.2. Plant phenotyping

Plant phenotyping/phenomics in crop breeding is to quantify various plant phenotypes (e.g., growth dynamics, stress resistance) resulting from the interactions of genetics and environmental conditions. Imaging technologies play a crucial role in the realization of high throughput, automated quantification of plant phenotypes (Minervini et al., 2015; Das Choudhury et al., 2019), thereby accelerating breeding processes and bridging the genotype-phenotype gap (Furbank and Tester, 2011). Imaging-based plant phenotyping is still facing challenges with extracting meaningful phenotypic information from images due to factors such as lighting variations, plant rotations, occlusions (Das Choudhury et al., 2019). Data-driven machine learning methods have been leveraged by plant scientists for effective feature extraction, plant trait identification/classification and quantification (Singh et al., 2018; Jiang and Li, 2020). To harness DL in plant phenotyping requires the curation of sufficient, high-quality labeled training samples, where GANs have emerged as powerful methods for synthetically expanding plant image datasets (Table 5).

Giuffrida et al. (2017) presented one of the earliest studies on image synthesis for plant phenotyping. The authors adapted DCGAN (Radford et al., 2015) with the number of leaves of a plant as the conditional input, trained on the CVPPP2017 dataset of potted *Arabidopsis* plants (Minervini et al., 2016), which was able to generate realistic color images of 128×128 pixels. Later, Zhu et al. (2018) reported on using Pix2Pix (Isola et al., 2017) to synthesize plant images of 512×512 pixels for leaf counting. The GAN model was trained on a collection of 500 images from the CVPPP 2017 dataset, along with corresponding leaf segmentation masks. The authors evaluated their approach by applying MaskRCNN (He et al., 2017) on the augmented dataset, obtaining 16.67% reduction in the leaf counting error compared to training the network on the original dataset without data augmentation. These two studies are however limited to generating samples for a single plant species. Madsen et al. (2019a) applied a Wasserstein auxiliary classifier GAN (WacGAN) that combined ACGAN (Odena et al., 2017) and WGAN-GP (Gulrajani et al., 2017), to generate plant seedling images of multiple species. The WacGAN yielded $58.9 \pm 9.2\%$ plant recognition accuracy on the 9-species plant seedling dataset (Giselsson et al., 2017), evaluated using ResNet-101 (He et al., 2016). Subsequently, based on WacGAN (Madsen et al., 2019a) and InfoGAN (Chen et al., 2016), Madsen et al. (2019b) presented WacGAN-info that incorporates unsupervised latent input variables in the GAN configuration for generating plant seedling images, achieving an improved classification accuracy of 64.3%. However, the GAN model produced low-resolution images of 128×128 pixels, and its impact on plant classification was not examined in comparison with training on either the original data without data augmentation or images augmented by traditional methods.

To deal with insufficient data in plant vigor rating, Zhu et al. (2020) applied conditional DCGAN (cDCGAN) to the images of orchid seedlings of two vigor ratings (healthy and weak). The *generator* and *discriminator* networks were conditioned by concatenating one-hot class labels at all

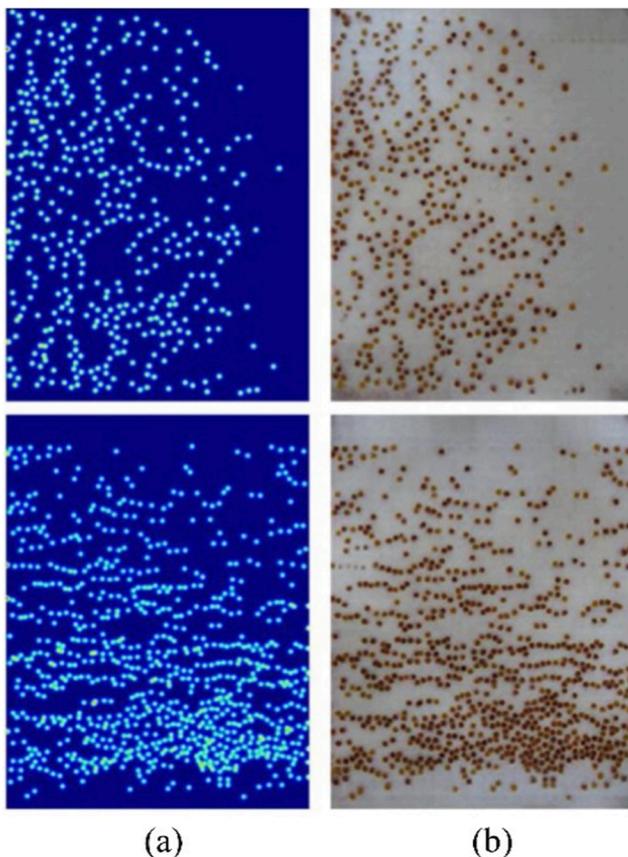


Fig. 17. Examples of synthetic shrimp egg images generated by CycleGAN: (a) randomly generated ground-truth density maps and (b) corresponding synthetic images. Reproduced with permission from Zhang et al. (2021a).

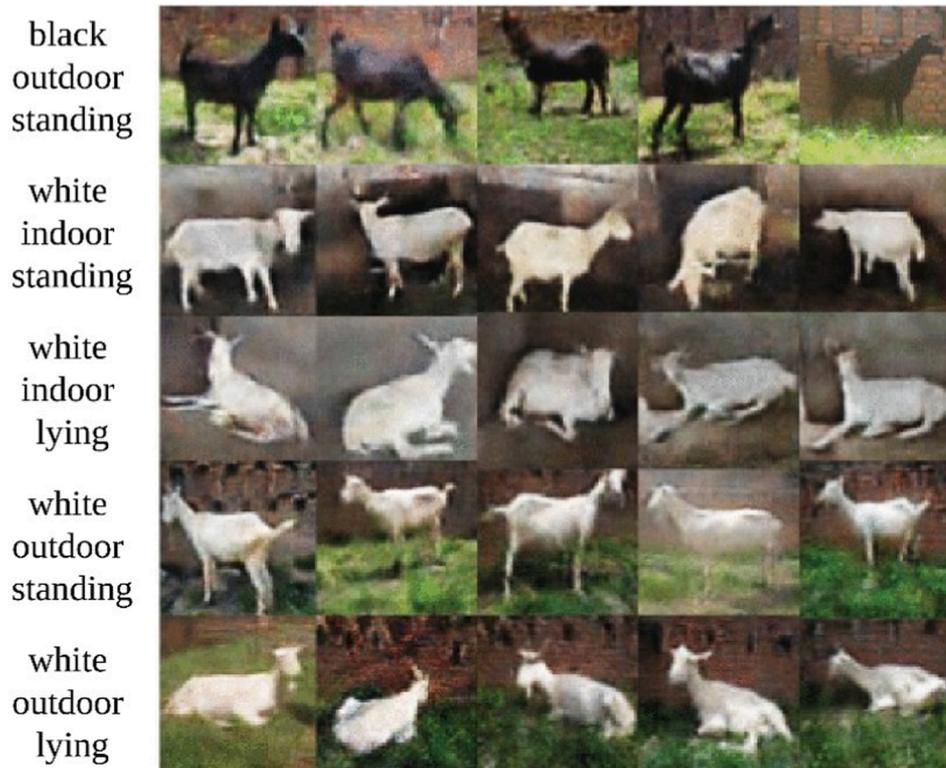


Fig. 18. Dairy goat images generated by improved SA-GAN with corresponding multi-labels (Li & Tang, 2020).

layers for image generation, and the classification using ResNet-50 (He et al., 2016) with the augmented data achieved a 23% increase in the testing F1 score compared to the model trained without data augmentation. The authors also showed that adding bypass deconvolutional connections in the GAN architecture helped generate more realistic images. Drees et al. (2021) assessed the utility of CGAN based on Pix2Pix (Isola et al., 2017) for predicting plant images at future growth stages (Fig. 19). The GAN model was evaluated on two open-source datasets, i. e., the Arabidopsis plant dataset collected indoors (Bell and Dee, 2016) and the field collected cauliflower dataset (Bender et al., 2020). For instance, segmentation of plants by Mask-RCNN (He et al., 2017), the correction between the size of predicted and reference images reached $R2 = 0.95$ for the cauliflower dataset but decreased to the values of 0.66–0.82 for the second dataset. Several factors influenced the image generation, such as geo-reference of images (required for Pix2Pix image alignment), plant position and overlaps of neighboring plants, which restrict the applicability of Pix2Pix-based GANs to complex field images.

Collecting large-scale datasets from diverse fields for plant phenotyping is a lengthy process and resource-intensive. Existing phenotyping datasets mostly consist of only images acquired under well-controlled laboratory conditions. To address the paucity of field-based data, Shete et al. (2020) proposed TasselGAN modified from DCGAN (Radford et al., 2015) to generate images of maize tassels against synthetic sky backgrounds. The authors generated maize tassel images from the indoors collected data (Gage et al., 2017) and synthetic sky background patches using the data in (Dev et al., 2016; Li et al., 2011) separately, and then merged them to form field-like images. However, the perceptual quality of generated images needs improvement, given the fact less than 40% of the generated images were considered sufficiently realistic based on qualitative assessment. The resultant image resolution (128×128 pixels) is not well suitable for phenotyping complex plant traits. An opensource field dataset called Global Wheat-Head Dataset (GWHD) (David et al., 2020), comprising 4700 high resolution RGB images and 190K labeled wheat heads collected from 12 countries, was recently released to facilitate wheat phenotyping research, particularly on the

detection of wheat heads. The dataset, however, has domain or distribution shift issues (Koh et al., 2021; Zhu et al., 2018) with substantial variations across growth region and crop varieties, which can degrade the performance of machine learning models. While simulated or synthetic data can alleviate the issue (Najafian et al., 2021), a domain gap exists between the simulated and read images. To tackle domain shifts, Hartley and French (2021) used CycleGAN (Zhu et al., 2017) to perform synthetic to real domain adaptation for the GWHD dataset (David et al., 2020). The GAN model was trained using a synthetic wheat head dataset, containing over 5000 images with over 100K annotations, and part of the GWHD data as source and target domains, respectively. The wheat head detection was done by Mask-RCNN via Detectron2 (Wu et al., 2019) and mIoU gains of about 3–4% were achieved by using the combination of synthetic and real samples compared to baseline models using real wheat head data alone. Only using synthetic data, however, leads to dramatically deteriorated accuracy, highlighting the value of labeled real images.

5.3. Postharvest quality assessment

Machine vision or imaging technologies are pervasive for post-harvest quality assessment of agricultural products (Chen et al., 2002; Blasco et al., 2017; Lu et al., 2020). At modern packing facilities, machine vision (color or monochromatic imaging) systems have been fairly well established for automated grading and sorting of food products for quality such as shape, size, and color. Because of significant biological variations, tissue defects are however still manually inspected for many horticultural products. DL algorithms hold promise to empower machine vision systems with enhanced capabilities for defect detection of agricultural produce. Several studies (Table 5) have been performed recently on the use of GANs for synthesizing fruit images for defect detection described below.

Chou et al. (2019) used GAN in a DL pipeline for identifying defective coffee beans. The defect inspection module was comprised of YOLOv3 (Redmon and Farhadi, 2018) and Hough circle detectors combined with

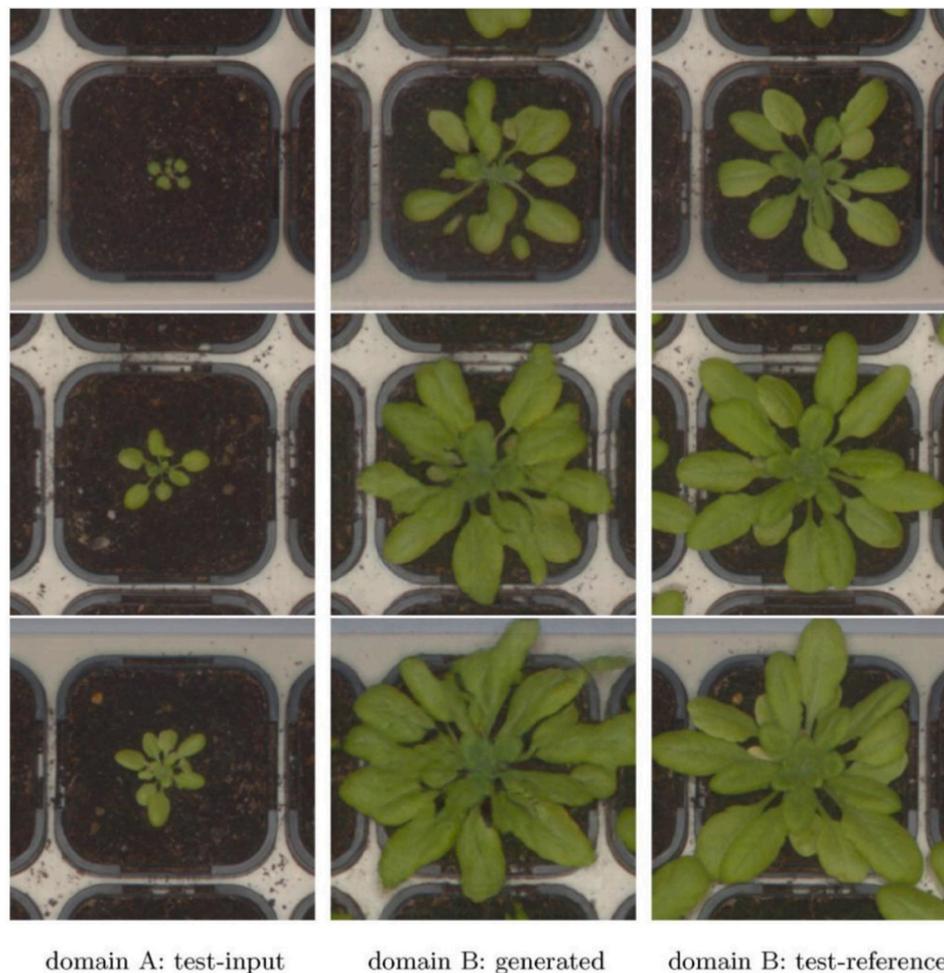


Fig. 19. Generated plant images at future growth stages with CGAN through domain adaptation, transferring early growth stages (domain A) to future growth stages (domain B). Images in the right column are those with the same domain as the generated image. Reproduced with permission from Drees et al. (2021)

a GAN-structured image generation mechanism, which could achieve detection rates of up to 80% for defective beans. However, the pixel resolution of generated images is still limited in this study, which can be detrimental to tasks like segmentation and quantification of defect areas. Guo et al. (2021) used DCGAN to generate jujube images for quality grading (1 healthy and 6 defective classes). Images (256×256 pixels) for defective classes were synthesized to alleviate the imbalance between healthy and defective samples. The inclusion of synthetic samples by DCGAN achieved classification accuracies of up to 98%, representing performance gains of 2–4% compared to models based on traditional data augmentation. It is noted that the GAN images for each defect class were generated separately rather than through a multi-class data generation process, which can be potentially more efficient. Recently, Bird et al. (2022) applied CGAN to generate synthetic lemon images for classifying healthy and unhealthy fruit, using a public lemon dataset (Adamiak, 2020). By using synthetic fruit images, the classification based on VGG16 achieved an accuracy of 88.75% against 83.77% without data augmentation. With a varied number (200 – 3000) of synthetic images for model training, the best accuracy was attained by using only 400 images (~13.5% of the whole dataset). The authors did not examine reasons for this phenomenon, which may be due to the quality of generated images or the capacity of classification models.

In real-world applications, imaging systems may be subjected to modifications (e.g., lighting conditions) over time, which can lead to dataset shift issues (Quinero-Candela et al., 2008). To address dataset shift problems in potato quality control, Marino et al. (2020) reported an

unsupervised adversarial domain adaptation (ADA) technique, which integrates adversarial learning and domain adaptation similarly to GNAs, for potato defect classification. Two datasets were collected for white potatoes and red tubers of six classes (1 healthy and 5 defective), respectively. The authors conducted two domain adaptation tasks, i.e., simulating lighting condition changes on the target domain by artificially increasing the brightness of part of white potato images and translating white potato to red tuber images. In the framework of adversarial discriminative domain adaptation (Tzeng et al., 2017), their ADA method with GoogleNet as the target classifier achieved significant performance gains on the target domain, with an average F1-score increasing to 0.84 from 0.46 in case of no domain adaptation. These approaches could be beneficial for enhancing the performance of imaging-based quality inspection (e.g., automated sorting) of horticultural products.

Transformer-based GANs have been proposed recently for improving the state of the art in image generation (Jiang et al., 2021; Zhao et al., 2021a; Hudson and Zitnick, 2021). Wang and Xiao (2021) reported on using TransGAN (Section 4.11) to generate lychee images for surface defect detection. Images were preferentially generated for two defect classes to rebalance training data, and the augmented-data achieved improvements of 0.58%–2.86% in mAP for defect detection, depending on object detectors, compared to training without GAN-based data augmentation. It is noted that the authors included only part of generated images for model training, because TransGAN did not always generate high-quality images. GANs have also been employed for

generating non-imaging sensor data for quality assessment of agricultural products. In Yang et al. (2021), 2D spectrograms of impact acoustic signals collected with a microphone were used for classifying wheat kernels (two defect classes along with healthy kernels). The authors used a vanilla GAN to generate synthetic spectrograms for data augmentation, obtaining an F1 score of 96.2% by using a PANSNet-5 network (Liu et al., 2018). This study however did not examine performance gains due to the GAN-generated data, which would require ablation studies without using GAN samples.

6. Challenges of GANs

Although GANs have demonstrated impressive performance in generating realistic images and improving DL models in agricultural applications, there remain several open challenges with training and evaluation, as identified and discussed below, alongside potential opportunities, which need to be well understood for exploiting the benefits of GANs in agriculture.

6.1. Training instability

Training GANs is prone significantly to instability problems (Arjovsky and Bottou, 2017; Goodfellow et al., 2020), despite a set of regularization techniques proposed for improving training stability (e.g., model architectures, loss function modifications, gradient penalty, spectral normalization) (Wiatrak et al., 2019; Liu et al., 2021). Two common failures of training include vanishing gradients (Goodfellow, 2017), leading to convergence failures, and mode collapse or dropping (Goodfellow, 2017; Durall et al., 2020; Srivastava et al., 2017), where the generator fails to learn well the distribution of training datasets.

GAN solves the minmax game [as formulated by Eq. (1)] by achieving the Nash equilibrium through a gradient descent method. When the generator is not as good as the discriminator, the discriminator always differentiates between generated and real samples and the gradients of the generator will vanish (close to 0), slowing or completely stopping the learning. On the other hand, mode collapse often happens, appearing as that the training data cannot be generated by the generator. This is probably the biggest issue with GANs (Goodfellow et al., 2017). The generator may produce only few modes of training samples (partial

mode collapse), and in the worst cases, simply the same set of instances (complete mode collapse) (Adler and Lunz, 2018; Arjovsky et al., 2017; Arjovsky and Bottou, 2017), as illustrated in Fig. 20. Mode collapse restricts the ability of GAN to generate diverse images and is very detrimental for visual recognition tasks.

Currently even the state-of-the-art GANs (e.g., BigGAN and TransGAN) would suffer from instability issues when trained on complex datasets and cannot be guaranteed to consistently produce images with high enough quality (Wang and Xiao, 2021; Yuwana et al., 2020; Espejo-Garcia et al., 2021; Shete et al., 2020). Moreover, the performance of GANs is generally sensitive to the settings of hyper-parameters (Kurach et al., 2019; Espejo-Garcia et al., 2021). Hence to make GANs work often requires significant manual efforts on hyperparameter tuning, network structure engineering, and a number of non-trivial training tricks (Kurach et al., 2019; Lucic et al., 2019). This is conceivably true for agricultural applications where GANs are used for modeling the distribution of highly complex, multi-class/-modal datasets, especially for those collected from unstructured, cluttered field conditions. Additionally, imbalanced training data, which are frequently encountered in agricultural applications, can contribute to GAN training difficulty and mode collapse (Dieng et al., 2019; Douzas and Bacao, 2018; Mariani et al., 2018).

6.2. GAN evaluation

Performance evaluation of GANs is another daunting task (Goodfellow, 2017). Currently there are no standardized, universal metrics for evaluating GANs and the quality of generated samples. Conventional error metrics (e.g., root mean square error) are not suitable because of a lack of direct one-to-one correspondence between generated and real images. Qualitative, visual assessment of the perceptual quality of generated images is commonly performed in computer vision literature (Ledig et al., 2017; Zhou et al., 2019), but it is subjective, prone to human error, time-consuming for large datasets, and may not capture distributional characteristics of datasets. Borji (2019, 2022) reviewed quantitative metrics for GAN evaluation, among which Inception Score (IS) (Salimans et al., 2016), F1-score, precision and recall (Lucic et al., 2019), Fréchet Inception Distance (FID) (Heusel et al., 2017), mode score (Che et al., 2016) and the metrics for image structural similarity

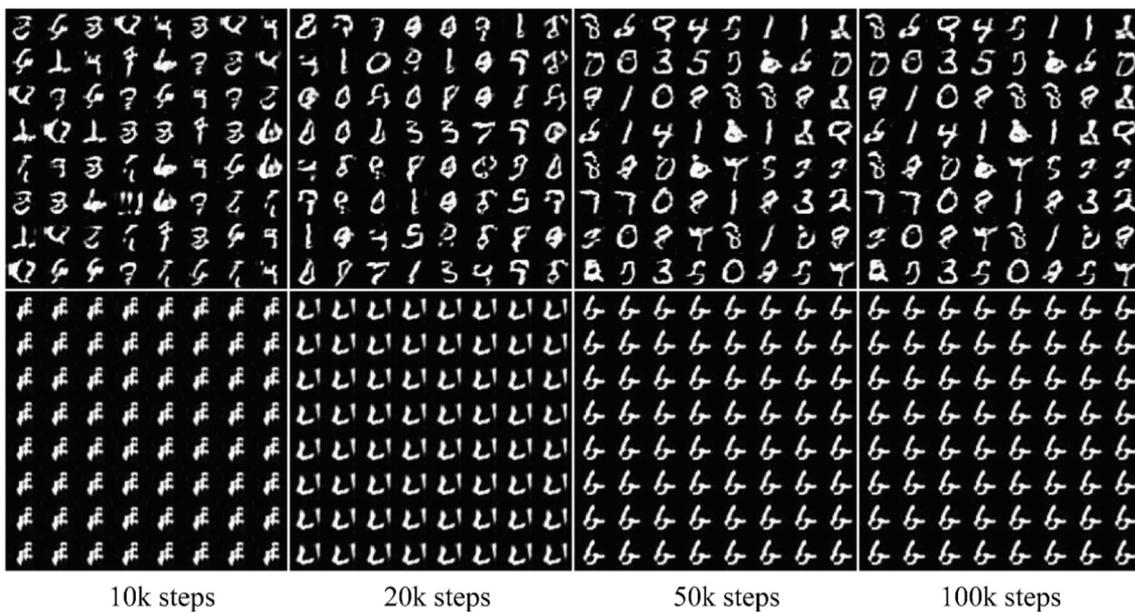


Fig. 20. Example of the model collapse (i.e., complete collapse) problem during GAN training on MNIST. The upper row shows the normal training procedure without model collapse. The bottom row shows an abnormal training with model collapse, where the generator only learns to output the digit “6”. Reproduced with permission from Metz et al. (2016).

assessment (Wang et al., 2004) are presumably the most popular for assessing the similarity between real and generated images. All of them, however, have shortcomings (Xu et al., 2018; Borji, 2019, 2022). There are also other metrics proposed for assessing the fidelity and diversity characteristics of generated images (Naeem et al., 2020; Kynkaanniemi et al., 2019). For agricultural applications reviewed in Section 5, only a handful of studies (Nazki et al., 2020; Shete et al., 2020; Espejo-Garcia et al., 2021) reported on quantitative metrics for assessing GAN-generated images. The validity of these metrics for agricultural images remains to be explored at large.

Instead of quantitatively evaluating the quality of generated images, most applications in Section 5 examined the impact of GANs on the eventual performance of DL models in visual recognition tasks, in comparison to baseline models using traditional image augmentation or the original, un-augmented data. These studies provide overwhelming evidence that DL models can substantially benefit from GAN-augmented data. For instance, Xu et al. (2022) observed the plant disease recognition rate increasing from 86.75% of the original data to 94.61% based on the ensemble of the original, GAN synthesized images and the data from traditional augmentation. Abbas et al. (2021) obtained 97.11% accuracy on GAN-augmented images compared to 94.34% on original. Sun et al. (2020) reported performance gains of about 3% due to GAN for plant disease classification. A more substantial performance gain of about 10% was achieved in Zhu et al. (2020) on plant vigor rating.

Although it seems to be easier to evaluate GANs on downstream machine learning tasks, quantitative evaluation of the quality and diversity of images generated by GANs is important for gaining an in-depth understanding of their strengths and limitations and inspire GAN innovations and hence should be explored in future research.

6.3. Training with limited data

High-quality image generation by GANs (Brock et al., 2019; Donahue and Simonyan, 2019; Karras et al., 2019, 2020b) is fueled by the seemingly sufficient supply of training images. For agricultural applications (e.g., plant disease detection) where there are only a small number of training images available, it can be difficult to train GANs to generate meaningful images beneficial for downstream DL model development. For instance, in Zhu et al. (2020) where GAN models were trained with varied numbers of training images to generate orchid seedlings, the generated images were perceptually inferior with the smallest number of training images, failing to capture detailed structure on leaves and roots. Compared to real samples, low-quality generated images may lack texture details and contain unrealistic, undesirable artifacts. In Bird et al. (2022) for generating lemon images, many synthetic images were found more reminiscent of potatoes than lemons and some suffered from unrealistic checkboard patterns. Provision of sufficient samples for training GANs would facilitate generating high-quality, realistic images, and eventually benefit DL models.

The key issue with limited data for training GANs is fundamentally that the *discriminator* overfits to training samples and consequently its feedback to the *generator* become useless, leading to the divergence of the training process (Salimans et al., 2016; Arjovsky et al., 2017). The traditional image augmentation methods described in Section 3, can be potentially used to prevent the *discriminator* from overfitting and facilitate GAN training. In Bi and Hu (2020) on plant disease classification, traditional methods (e.g., flipping, rotation and contrast manipulation) were first used to augment the original data, followed by GAN training to generate synthetic images to train CNN models. This strategy was proven effective for achieving high image classification accuracies, which were also used in other studies (Arsenovic et al., 2019; Douarre et al., 2019).

In addition, to train GANs with limited data, Karras et al. (2020a) proposed an adaptive *discriminator* augmentation mechanism that reliably stabilizes training and improves image generation of GANs when training data is in short supply. The authors also acknowledged that the

augmentation should not be substituted for real data to train high-quality GANs. Tran et al. (2021) recognized the issues with classic data augmentation for GANs and proposed a principled framework of data augmentation optimized to improve the learning of *discriminator* and *generator*. These methods can be potentially useful for helping train GANs with limited data in agricultural applications.

6.4. Other considerations

As discussed above, GAN models have the risk of generating low-quality images, while this may not be a real problem for traditional image augmentation methods that generally do not significantly impair image quality. That said, GAN-based image augmentation may not always translate into performance gains of DL models, compared to traditional image augmentation or using the original data, which is observed in many application studies (Xu et al., 2022; Hartley and French, 2021; Douarre et al., 2019; Fawakherji et al., 2021). However, strong evidence shows consistent improvements can be obtained by using both GAN and traditional augmentation than performing them separately, or by mixing the original and GAN-augmented data. This suggests that GAN and traditional image augmentation can work synergistically to improve the performance of DL models. More extensive research is still needed on systematic evaluation of the impacts of GANs and traditional image augmentation in agricultural applications.

Semi-supervised GANs aim to learn from data in a semi-supervised fashion, where a small number of labeled images with large unlabeled data is to achieve model performance comparable to supervised training. Kerdegari et al. (2019) successfully applied semi-supervised GAN to weed images with only half of labeled training data. Semi-supervised GANs should be well considered in case of the difficulty of labelling large-scale data. Makeup removal (removing irrelevant backgrounds or occlusions) is found effective for assisting in image generation of GANs. Kierdorf et al. (2021) applied Pix2Pix GAN to remove the leaves from the berry images to generate non-occluded berry images for accurate berry counting. Olatunji et al. (2020) reconstructs the complete kiwifruit surface by translating occluded fruit into non-occluded images to address partial occlusion problems in fruit detection. More research on image processing operations like makeup removal can be beneficial to the application of GANs to visual recognition tasks such as fruit detection and counting.

Multi-task learning is a modeling paradigm that enables simultaneously learning multiple related tasks with the goal to improve the generalization power for all the tasks (Zhang and Yang, 2021; Vandenhende et al., 2021). Currently, most GAN applications in agriculture focus on generating images for a single task (e.g., image classification or object detection). Training in a multi-task manner can potentially foster generalization of GANs and improve overall performance for each task by leveraging knowledge contained in other tasks (Zhang and Yang, 2021; Vandenhende et al., 2021). A multi-task GAN model can be designed by using a multi-task discriminator network. For instance, in Bai et al. (2018), an end-to-end multi-task GAN was proposed for detecting small objects, in which the discriminator performs multiple tasks including distinguishing real from generated images, predicting object categories, and regressing bounding boxes simultaneously. This model showed superior object detection performance on the COCO dataset (Lin et al., 2014). Currently there is scant research on multi-task GANs in agriculture, which would be worthy of future investigation.

7. Summary

Since its first proposal in 2014, GAN has attracted growing interest in computer vision areas. The ability of GANs to synthesize plausible, diverse images opens new opportunities to enable better-performing DL models trained for agricultural applications, especially when large-scale, labeled image datasets are not readily available. This paper makes a first comprehensive review of GANs in agricultural contexts. We have

presented an overview of traditional image augmentation methods and the landscape of GAN architectural variants. Through a systematic literature review, we have surveyed a total of 59 publications of GANs in agriculture since 2017, for image augmentation in precision agriculture, plant phenotyping and postharvest inspection of agricultural produce. GANs have achieved remarkable performance in a diversity of visual recognition tasks in agriculture and food systems, such as plant disease classification, weed detection, fruit counting, postharvest fruit defect detection, etc. We also discuss challenges and opportunities regarding GAN training and evaluation. This paper will be beneficial for facilitating further research and application of GANs in agriculture and food domains.

Credit authorship contribution statement

Yuzhen Lu: Conceptualization, Investigation, Supervision, Writing – original draft & review & editing. **Dong Chen:** Formal analysis, Software, Writing – original draft. **Ebenezer Olaniyi:** Formal analysis, Software, Writing – original draft. **Yanbo Huang:** Writing – review & editing.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

This work was supported in part by Cotton Incorporated award #21-005 and the USDA National Institute of Food and Agriculture Hatch project #1025922.

References

- Abadi, M., Barham, P., Chen, J., Chen, Z., Davis, A., Dean, J., Devin, M., Ghemawat, S., Irving, G., Isard, M., Kudler, M., Levenberg, J., Monga, R., Moore, S., Murray, D.G., Steiner, B., Tucker, P., Vasudevan, V., Warden, P., Wicke, M., Yu, Y., Zheng, X., 2016. TensorFlow: A system for large scale machine learning. In: The 12th USENIX Symposium on Operating Systems Design and Implementation, pp. 265–283.
- Abbas, A., Jain, S., Gour, M., Vankudothu, S., 2021. Tomato plant disease detection using transfer learning with CGAN synthetic images. *Comput. Electron. Agric.* 187, 106279.
- Abdoli, A., Murillo, A.C., Yeh, C.C.M., Gerry, A.C., Keogh, E.J., 2018. Time series classification to improve poultry welfare. In: 17th IEEE International conference on machine learning and applications, pp. 635–642.
- Adamiak, M., 2020. Lemons quality control dataset, <https://github.com/softwaremill/lemon-dataset>.
- Adler, J., Lunz, S., 2018. Banach Wasserstein generative adversarial network. In: The 32nd Conference on Neural Information Processing Systems, pp. 6755–6764.
- Affif, M., Price, B., Cohen, S., Brown, M. S., 2019. When color constancy goes wrong: correcting improperly white-balanced images. In: proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 1535–1544.
- Ahmed, G., Malick, R. A. S., Akhuzada, A., Zahid, S., Sagriand, M. R., Gani, A., 2021. An approach towards IoT-based predictive service for early detection of diseases in poultry chickens. *Sustainability* 13 (23), 13396.
- Arjovsky, M., Bottou, L., 2017. Towards principled methods for training generative adversarial networks. In: 5th International Conference on Learning Representations (ICLR), pp. 1–17.
- Arjovsky, M., Chintala, S., Bottou, L., 2017. Wasserstein generative adversarial networks. In: Proceedings of the 34th International Conference on Machine Learning 70, 214–223.
- Arsenovic, M., Karanovic, M., Sladojevic, M., Anderla, A., Stefanovic, D., 2019. Solving current limitations of deep learning-based approaches for plant disease detection. *Symmetry* 11 (7), 939.
- Badrinarayanan, V., Kendall, A., Cipolla, R., 2017. Segnet: A deep convolutional encoder-decoder architecture for image segmentation. *IEEE Trans. Pattern Anal. Mach. Intell.* 39 (12), 2481–2495.
- Bai, Y., Zhang, Y., Ding, M., Ghanem, B., 2018. Sod-mtgan: Small object detection via multi-task generative adversarial network. In: Proceedings of the European Conference on Computer Vision (ECCV), 1–16.
- Bailey, B.N., 2019. Helios: a scalable 3D plant and environmental biophysical modeling framework. *Front. Plant Sci.* 10, 1185.
- Barth, R., Hemming, J., Van, E., Henten, J., 2018. Improved part segmentation performance by optimizing realism of synthetic images using cycle generative adversarial networks. arXiv preprint arXiv:1803.06301.
- Barth, R., Hemming, J., Van, E., Henten, J., 2020. Optimising realism of synthetic images using cycle generative adversarial networks for improved part segmentation. *Comput. Electron. Agric.* 173, 105378.
- Bau, D., Zhu, J.Y., Strobel, H., Zhou, B., Tenenbaum, J.B., Freeman, W.T., Torralba, A., 2018. Gan dissection: Visualizing and understanding generative adversarial networks. arXiv preprint arXiv:1811.10597.
- Bechar, A., Vigneault, C., 2016. Agricultural robots for field operations: Concepts and components. *Biosystems Engineering* 149, 94–111.
- Bell, J., Dee, H.M., 2016. Aberystwyth leaf evaluation dataset, <https://zenodo.org/record/168158#YldIC8hKi3A>.
- Bellocchio, E., Costante, G., Cascianelli, S., Fravolini, M.L., Valigi, P., 2020. Combining domain adaptation and spatial consistency for unseen fruits counting: a quasi-supervised approach. *IEEE Rob. Autom. Lett.* 5 (2), 1079–1086.
- Bender, A., Whelan, B., Sukkariéh, S., 2020. A high-resolution, multimodal data set for agricultural robotics: A ladybird's-eye view of brassica. *J. Field Rob.* 37 (1), 73–96.
- Bengio, Y., Lecun, Y., Hinton, G., 2021. Deep learning for AI. *Commun. ACM* 64 (7), 58–65.
- Berthelot, D., Schumm, T., Metz, L., 2017. Began: Boundary equilibrium generative adversarial network arXiv preprint arXiv:1703.10717.
- Bi, L., Hu, G., 2020. Improving image-based plant disease classification with generative adversarial network under limited training set. *Front. Plant Sci.* 11, 583438.
- Bird, J. J., Barnes, C. M., Manso, L. J., Ekart, A., D. R. Faria, D. R., 2022. Fruit quality and defect image classification with conditional GAN data augmentation. *Scientia Horticulturae* 293 (5), 1–11.
- Bissoto, A., Valle, E., Avila, S., 2021. Gan-based data augmentation and anonymization for skin-lesion analysis: A critical review. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 1847–1856.
- Blasco, J., Munera, S., Aleixos, N., Cubero, S., Molto, E., 2017. Machine vision-based measurement systems for fruit and vegetable quality control in postharvest. In: Hitzmann, B. (Eds.), *Measurement, Modeling and Automation in Advanced Food Processing. Advances in Biochemical Engineering/Biotechnology*, vol 161. Springer, Cham, 71–91.
- Bloice, M.D., Roth, P.M., Holzinger, A., 2019. Biomedical image augmentation using augmentator. *Bioinformatics* 35 (21), 4522–4524.
- Boom, A.J., Huang, P.X., He, J., Fisher, R.B., 2012. Supporting ground truth annotation of image datasets using clustering. In: Proceedings of the 21st International Conference on Pattern Recognition, 1542–1545.
- Borji, A., 2019. Pros and cons of GAN evaluation measures. *Comput. Vis. Image Underst.* 179, 41–65.
- Borji, A., 2022. Pros and cons of GAN evaluation measures: new developments. *Comput. Vis. Image Underst.* 215, 103329.
- Bowles, C., Chen, L., Guerrero, R., Bentley, P., Gunn, R., Hammers, A., Dickie, D.A., Hernandez, M. V., Wardlaw, J., Rueckert, D., 2018. GAN augmentation: augmenting training data using generative adversarial networks. arXiv preprint arXiv:1810.10863.
- Brahimi, M., Boukhalfa, K., Moussaoui, A., 2017. Deep learning for tomato diseases: classification and symptoms visualization. *Appl. Artificial Intell.* 31(4), 299–315.
- Brock, A., Donahue, J., Simonyan, K., 2019. Large scale GAN training for high fidelity natural image synthesis. In: 7th International Conference on Learning Representations (ICLR), pp. 1–34.
- Buslaev, A., Iglovikov, V.I., Khvedchenya, E., Parinov, A., Druzhinin, M., Kalinin, A.A., 2020. Albumentations: fast and flexible image augmentations. *Information* 11 (2), 125.
- Cao, Y.J., Jia, L.L., Chen, Y.X., Lin, N., Yang, C., Zhang, B., Liu, Z., Li, X.X., Dai, H.H., 2018. Recent advances of generative adversarial networks in computer vision. *IEEE Access* 7, 14985–15006.
- Cap, Q.H., Uga, H., Kagiwada, S., Iyatomi, H., 2020. Leafgan: An effective data augmentation method for practical plant disease diagnosis. *IEEE Trans. Autom. Sci. Eng.* 19 (2), 1258–1267.
- Carion, N., Massa, F., Synnaeve, G., Usunier, N., Kirillov, A., Zagoruyko, S., 2020. End-to-end object detection with transformers. In: European Conference on Computer Vision, 213–229.
- Che, T., Li, Y., Jacob, A.P., Bengio, Y., Li, W., 2016. Mode regularized generative adversarial networks. ArXiv preprint arXiv:1612.02136, 2016.
- Chebroly, N., Lottes, P., Schaefer, A., Winterhalter, W., Burgard, W., Stachniss, C., 2017. Agricultural robot dataset for plant classification, localization and mapping on sugar beet fields. *Int. J. Robot. Res.* 36 (10), 1045–1052.
- Chen, Y., Chao, K., Kim, M.S., 2002. Machine vision technology for agricultural applications. *Comput. Electron. Agric.* 36 (2), 173–191.
- Chen, X., Duan, Y., Houthoof, R., Schulman, J., Sutskever, I., Abbeel, P., 2016. Infogan: Interpretable representation learning by information maximizing generative adversarial nets. In: The 30th Conference on Neural Information Processing Systems, pp. 2180–2188.
- Chen, D., Lu, Y., Li, Z., Young, S., 2022. Performance evaluation of deep transfer learning on multi-class identification of common weed species in cotton production systems. *Comput. Electron. Agric.* 198, 107091.
- Chou, Y.C., Kuo, C.J., Chen, T.T., Horng, G.J., Pai, M.Y., Wu, M.E., Lin, Y.C., Hung, M.H., Su, W.T., Chen, Y.C., Wang, D.C., Chen, C.C., 2019. Deep-learning based defective bean inspection with GAN-structured automated labeled data augmentation in coffee industry. *Appl. Sci.* 9 (19), 4166.
- Ciresan, D.C., Meier, U., Gambardella, L.M., Schmidhuber, J., 2010. Deep, big, simple neural nets for handwritten digit recognition. *Neural Comput.* 22 (12), 3207–3220.

- Creswell, A., White, T., Dumoulin, V., Arulkumar, K., Sengupta, B., Bharath, A.A., 2018. Generative adversarial networks: An overview. *IEEE Signal Process Mag.* 35 (1), 53–65.
- Cubuk, E.D., Zoph, B., Mane, D., Vasudevan, V., Le, Q.V., 2019. Autoaugment: Learning augmentation strategies from data. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 113–123.
- Cui, S., Wei, M., Liu, C., Jiang, J., 2021. GAN-SegNet: A deep generative adversarial segmentation network for brain tumor semantic segmentation. *Int. J. Imaging Syst. Technol.* <https://doi.org/10.1002/ima.22677>.
- Dai, Q., Cheng, X., Qiao, Y., Zhang, Y., 2020a. Agricultural pest super-resolution and identification with attention enhanced residual and dense fusion generative and adversarial network. *IEEE Access* 8, 81943–81959.
- Dai, Q., Cheng, X., Qiao, Y., Zhang, Y., 2020b. Crop leaf disease image super-resolution and identification with dual attention and topology fusion generative adversarial network. *IEEE Access* 8, 55724–55735.
- Das Choudhury, S., Samal, A., Awada, T., 2019. Leveraging image analysis for high-throughput plant phenotyping. *Front. Plant Sci.* 10, 508.
- David, E., Maded, S., Sadeghi-Tehrani, P., Aasen, H., Zheng, B., Liu, S., Kirchgessner, N., Ishikawa, G., Nagasawa, K., Badhon, M.A., Pozniak, C., de Solan, B., Hund, A., Chapman, S.C., Baret, F., Stavness, I., Guo, W., 2020. Global wheat head detection (GWHD) dataset: a large and diverse dataset of high-resolution RGB-labelled images to develop and benchmark wheat head detection methods. *Plant Phenomics* 2020, 3521852.
- Davies, E., 2009. The application of machine vision to food and agriculture: a review. *Imaging Sci. J.* 57 (4), 197–217.
- Deng, J., Dong, Socher, R., Li, L.J., Li, K., Fei-Fei, L., 2009. Imagenet: A large-scale hierarchical image database. In: *IEEE Conference on Computer Vision and Pattern Recognition*, 248–255.
- Deng, H., Luo, D., Chang, Z., Li, H., Yang, X., 2021. Rahc GAN: A data augmentation method for tomato leaf disease recognition. *Symmetry* 13 (9), 1597.
- Dev, S., Lee, Y.H., Winkler, S., 2016. Color-based segmentation of sky/cloud images from ground based cameras. *IEEE J. Sel. Top. Appl. Earth Obs. Remote Sens.* 10 (1), 231–242.
- Dhaka, V.S., Meena, S.V., Rani, G., Sinwar, D., Ijaz, M.F., Wozniak, M., 2021. A survey of deep convolutional neural networks applied for prediction of plant leaf diseases. *Sensors* 21 (14), 4749.
- Dieng, A.B., Ruiz, F.J., Blei, D.M., Titsias, M.K., 2019. Prescribed generative adversarial networks. *arXiv preprint arXiv:1910.04302*.
- Donahue, J., Simonyan, K., 2019. Large scale adversarial representation learning. *arXiv preprint arXiv:1907.02544*.
- Douarre, C., Crispim-Junior, C.F., Gelibert, A., Tougne, L., Rousseau, D., 2019. Novel data augmentation strategies to boost supervised segmentation of plant disease. *Comput. Electron. Agric.* 165, 104967.
- Douzas, G., Bacao, F., 2018. Effective data generation for imbalanced learning using conditional generative adversarial networks. *Expert Syst. Appl.* 91, 464–471.
- Drees, L., Junker-Frohn, L.V., Kierdorf, J., Roscher, R., 2021. Temporal prediction and evaluation of brassica growth in the field using conditional generative adversarial networks. *Comput. Electron. Agric.* 190, 106415.
- Dumoulin, V., Shlens, J., Kudlur, M., 2016. A learned representation for artistic style. *arXiv preprint arXiv:1610.07629*.
- Durall, R., Chatzimichailidis, A., Labus, P., Keuper, J., 2020. Combating mode collapse in GAN training: An empirical analysis using hessian eigenvalues. *arXiv preprint arXiv:2012.09673*.
- Eppenhof, K.A., Lafarge, M.W., Veta, M., Pluim, J.P., 2019. Progressively trained convolutional neural networks for deformable image registration. *IEEE Trans. Med. Imaging* 39(5), 1594–1604.
- Espejo-Garcia, B., Mylonas, N., Athanasakos, L., Vali, E., Fountas, S., 2021. Combining generative adversarial networks and agricultural transfer learning for weeds identification. *Biosyst. Eng.* 204, 79–89.
- Fawakherji, M., Potena, C., Pretto, A., Bloisi, D.D., Nardi, D., 2021. Multi-Spectral Image Synthesis for Crop/Weed Segmentation in Precision Farming. *Rob. Auton. Syst.* 146, 103861.
- Fei, Z., Olenskyj, A., Bailey, B.N., Earles, M., 2021. Enlisting 3D crop models and GANs for more data efficient and generalizable fruit detection. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 1269–1277.
- Fore, M., Frank, K., Norton, T., Svendsen, E., Alfredsen, J.A., Dempster, T., Eguiraun, H., Watson, W.A., Stahl, A., Sunde, L.M., 2018. Precision fish farming: A new framework to improve production in aquaculture. *Biosyst. Eng.* 173, 176–193.
- Frid-Adar, M., Diamant, I., Klang, E., Amitai, M., Goldberger, J., Greenspan, H., 2018. GAN-based synthetic medical image augmentation for increased CNN performance in liver lesion classification. *Neurocomputing* 321, 321–331.
- Furbank, R.T., Tester, M., 2011. Phenomics—technologies to relieve the phenotyping bottleneck. *Trends Plant Sci.* 16 (12), 635–644.
- Gage, J.L., Miller, N.D., Spalding, P., Kaepler, S.M., de Leon, N., 2017. Tips: a system for automated image-based phenotyping of maize tassels. *Plant Methods* 13 (1), 1–12.
- Giselsson, T.M., Jørgensen, R.N., Jensen, P.K., Dyrmann, M., Midtby, H.S., 2017. A public image database for benchmark of plant seedling classification algorithms. *arXiv Preprint arXiv:1711.05458*.
- Giuffrida, M.V., Scharr, H., Tsafaris, S.A., 2017. ARIGAN: Synthetic arabidopsis plants using generative adversarial network. In: *IEEE International Conference on Computer Vision Workshops*, 2064–2071.
- Gomaa, A.A., El-Latif, Y.M.A., 2021. Early prediction of plant diseases using CNN and GANs. *Int. J. Adv. Comput. Sci. Appl.* 12 (5), 514–519.
- Gong, X., Chang, S., Jiang, Y., Wang, Z., 2019. Autogan: Neural architecture search for generative adversarial networks. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 3224–3234.
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., 2014. Generative adversarial nets. In: *Advances in Neural Information Processing Systems*, 2672–2680.
- Goodfellow, I., Bengio, Y., Courville, A., Bengio, Y., 2016. *Deep Learning*. MIT Press, Cambridge, MA.
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y., 2020. Generative adversarial networks. *Commun. ACM* 63 (11), 139–144.
- Goodfellow, I., 2017. Nips 2016 tutorial: Generative adversarial networks. *arXiv preprint arXiv:1701.00160*.
- Gui, J., Sun, Z., Wen, Y., Tao, D., Ye, J., 2020. A review on generative adversarial networks: Algorithms, theory, and applications. *J. Latex Class file* 14 (8), 1–20.
- Gulrajani, I., Ahmed, F., Arjovsky, M., Dumoulin, V., Courville, A., 2017. Improved training of wasserstein GANs. In: *Proceedings of the 31st International Conference on Neural Information Processing Systems*, 5768–5778.
- Guo, Z., Zheng, H., Xu, X., Ju, J., Zheng, Z., You, C., Gu, Y., 2021. Quality grading of jujubes using composite convolutional neural networks in combination with RGB color space segmentation and deep convolutional generative adversarial networks. *J. Food Process Eng.* 44 (2), e13620.
- Gupta, A., Dollar, P., Girshick, R., 2019. Lvis: A dataset for large vocabulary instance segmentation. In: *Proceeding IEEE Computer Society Conference Computer Vision Pattern Recognition*, 5351–5359.
- Han, K., Guo, J., Tang, Y., and Wang, Y. 2022. Pyramidtn: Improved transformer-linear transformer baseline with pyramid architecture. *arXiv preprint arXiv:2201.00978*.
- Hartley, Z.K., French, A.P., 2021. Domain adaptation of synthetic images for wheat head detection. *Plants* 10 (12), 2633.
- He, K., Zhang, X., Ren, S., Sun, J., 2016. Deep residual learning for image recognition. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778.
- He, K., Gkioxari, G., Dollar, P., Girshick, R., 2017. Mask R-CNN. In: *Proceedings of the IEEE international conference on computer vision*, pp. 2961–2969.
- Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., Hochreiter, S., 2017. GANS trained by a two time-scale update rule converge to a local Nash equilibrium. In: *Proceedings of the 31st International Conference on Neural Information Processing Systems*, pp. 6629–6640.
- Ho, J., Jain, A., Abbeel, P., 2020. Denoising diffusion probabilistic models. In: *Proceedings of the 34th International Conference on Neural Information Processing Systems* 574, 6840–6851.
- Hu, W.J., Xie, T.Y., Li, B., Du, Y.-X., Xiong, N.N., 2021. An edge intelligence-based generative data augmentation system for lot image recognition tasks. *J. Internet Technol.* 22(4), 765–778.
- Hu, G., Wu, H., Zhang, Y., Wan, M., 2019. A low shot learning method for tea leaf's disease identification. *Comput. Electron. Agric.* 163, 104852.
- Huang, X., Belongie, S., 2017. Arbitrary style transfer in real-time with adaptive instance normalization. In: *Proceedings of the IEEE international conference on computer vision*, pp. 1501–1510.
- Huang, H., Yu, P.S., Wang, C., 2018. An introduction to image synthesis with generative adversarial nets. *arXiv preprint arXiv:1803.04469*.
- Hudson, D. A., Zitnick, L., 2021. Generative adversarial transformers. In: *International Conference on Machine Learning* 139, 4487–4499.
- Hughes, D., Salathe, M., et al. 2015. An open access repository of images on plant health to enable the development of mobile disease diagnostics. *arXiv preprint arXiv:1511.08060*.
- Ioffe, S., Szegedy, C., 2015. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In: *Proceedings of the 32nd International Conference on International Conference on Machine Learning*, 37, 448–456.
- Isola, P., Zhu, J.Y., Zhou, T., Efros, A.A., 2017. Image-to-image translation with conditional adversarial networks. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 1125–1134.
- Jager, J., Simon, M., Denzler, J., Wolff, V., Fricke-Neudert, K., Kruschel, C., 2015. Croatian fish dataset: Fine-grained classification of fish species in their natural habitat. In: *British Machine Vision Conference* 2, 1–7.
- Jiang, Y., Chang, S., Wang, Z., 2021. TransGAN: Two pure transformers can make one strong gan, and that can scale up. In: *34th Advances in Neural Information Processing Systems. NeurIPS Proceedings*.
- Jiang, Y., Li, C., 2020. Convolutional neural networks for image-based high-throughput plant phenotyping: a review. *Plant Phenomics* 2020, 4152816.
- Jolicoeur-Martineau, A., 2018. The relativistic discriminator: a key element missing from standard GAN. *arXiv preprint arXiv:1807.00734*.
- Jungo, A., Scheidegger, O., Reyes, M., Balsiger, F., 2021. Pymia: A python package for data handling and evaluation in deep learning based medical image analysis. *Comput. Methods Programs Biomed.* 198, 105796.
- Junior, G. O., Schaeffer, L., Schenkel, F., Tiezzi, F., Baes, C. F., Potential effects of hormonal synchronized breeding on genetic evaluations of fertility traits in dairy cattle: A simulation study. *J. Dairy Sci.* 104(4), 4404–4412.
- Kamilaris, A., Prenafeta-Boldú, F.X., 2018. *Deep learning in agriculture: A survey*. *Comput. Electron. Agric.* 147, 70–90.
- Karras, T., Aila, T., Laine, S., Lehtinen, J. 2018. Progressive growing of GANs for improved quality, stability, and variation. In: *International Conference on Learning Representations*, pp. 1–26.
- Karras, T., Aittala, M., Hellsten, J., Laine, S., Lehtinen, J., Aila, T., 2020a. Training generative adversarial networks with limited data. In: *34th Conference on Neural Information Processing Systems* 33, 12104–12114.
- Karras, T., Aittala, M., Laine, S., Harkonen, E., Hellsten, J., Lehtinen, J., Aila, T., 2021. Alias-free generative adversarial networks. In: *35th Conference on Neural Information Processing Systems*.

- Karras, T., Laine, S., Aila, T., 2019. A style-based generator architecture for generative adversarial networks. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 4401–4410.
- Karras, T., Laine, S., Aittala, M., Hellsten, J., Lehtinen, J., and Aila, T., 2020b. Analyzing and improving the image quality of stylegan. In: *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 8110–8119.
- Kerdegari, H. Razaak, M. Argyriou, V., Remagnino, P., 2019. Semi-supervised GAN for classification of multispectral imagery acquired by UAVs. arXiv preprint arXiv: 1905.10920.
- Khalifa, N.E., Mohamed, K., Seyedali, L., 2022. A comprehensive survey of recent trends in deep learning for digital images augmentation. *Artif. Intell. Rev.* 55 (3), 2351–2377.
- Khan, S., Tufail, M., Khan, M.T., Khan, Z.A., Iqbal, J., Alam, M., 2021. A novel semi-supervised framework for UAV based crop/weed classification. *PLoS ONE* 16 (5), e0251008.
- Kierdorf, J., Weber, I., Kicherer, A., Zabawa, L., Drees, L., Roscher, R., 2021. Behind the leaves – Estimation of occluded grapevine berries with conditional generative adversarial networks. arXiv preprint arXiv:2105.10325.
- Kim, C., Lee, H., Jung, H., 2021. Fruit tree disease classification system using generative adversarial networks. *Int. J. Electr. Comput. Eng.* 11 (3), 2508–2515.
- Kirillov, A., Wu, Y., He, K., Girshick, R., 2020. Pointrend: Image segmentation as rendering. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 9799–9808.
- Koh, P.W., Sagawa, S., Marklund, H., Xie, S.M., Zhang, A., Balsubramani, W., Hu, Yasunaga, M., Phillips R. L., Gao, I., 2021. Wilds: A benchmark of in-the-wild distribution shifts. In: *Proceedings of the 38th International Conference on Machine Learning*, 5637–5664.
- Koivala, A., Walsh, K.B., Wang, Z., McCarthy, C., 2019. Deep learning method overview and review of use for fruit detection and yield estimation. *Comput. Electron. Agric.* 162, 219–234.
- Konidaris, F., Tagaris, T., Sdraka, M., Stafylopatis, A., 2019. Generative adversarial networks as an advanced data augmentation technique for MRI data. In: *Proceedings of the 14th International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications (VISIGRAPP)*, pp. 48–59.
- Krizhevsky, A., Hinton, G., 2009. Learning multiple layers of features from tiny images. *University of Tront*, pp. 2–58. Master's thesis.
- Krizhevsky, A., Sutskever, I., Hinton, G.E., 2012. ImageNet Classification with Deep Convolutional Neural Networks. *Adv. Neural Inform. Process. Syst.* 25 (2), 1–9.
- Kurach, K., Lucic, M., Zhai, X., Michalski, M., Gelly, S., 2019. A large-scale study on regularization and normalization in GANs. In: *International Conference on Machine Learning*, 3581–3590.
- Kurutach, T., Tamar, A., Yang, G., Russell, S.J., Abbeel, P., 2018. Learning plannable representations with causal infoGAN. In: *The 32nd Conference on Neural Information Processing Systems*, pp. 8747–8758.
- Kuznetsova, A., Rom, H., Alldrin, N., Uijlings, J., Krasin, I., Pont-Tuset, J. Kamali, S. Popov, S., Mallocci, Kolesnikov, M., Duerig, A., Duerig, T., Ferrari, V. 2020. The open images dataset v4: unified image classification, object detection, and visual relationship detection at scale. *Int. J. Comput. Vision* 128, 1956–1981.
- Kynkaanniemi, T., Karras, T., Laine, T., Lehtinen, J., Aila, J., 2019. Improved precision and recall metric for assessing generative models. In: *The 33rd Conference on Neural Information Processing Systems* 353, 3927–3936.
- Lan, L., You, L., Zhang, Z., Fan, Z., Chen, Y., Zhou, X., 2020. Generative adversarial networks and its applications in biomedical informatics. *Front. Public Health* 8, 164.
- LeCun, Y., Bengio, Y., Hinton, G., 2015. Deep learning. *Nature* 521 (7553), 436–444.
- Ledig, C., Theis, L., Huszar, F., Caballero, J., Cunningham, A., Acosta, A., Aitken, A., Tejani, J., Wang, Z., 2017. Photo-realistic single image super-resolution using a generative adversarial network. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 4681–4690.
- Li, Q., Lu, W., Yang, J., 2011. A hybrid thresholding algorithm for cloud detection on ground-based color images. *J. Atmos. Oceanic Technol.* 28 (10), 1286–1296.
- Li, D., Miao, Z., Peng, F., Wang, L., Hao, Y., Wang, Z., Chen, T., Li, H., Zheng, Y., 2021a. Automatic counting methods in aquaculture: a review. *J. World Aquacult. Soc.* 52 (2), 269–283.
- Li, H., Tang, J., 2020. Dairy goat image generation based on improved-self-attention generative adversarial networks. *IEEE Access* 8, 62448–62457.
- Li, J., Zhao, X., Zhou, G., Zhang, M., Li, D., Zhou, Y., 2021b. Evaluating the work productivity of assembling reinforcement through the objects detected by deep learning. *Sensors* 21 (160), 5598.
- Liakos, K.G., Busato, P., Moshou, D., Pearson, S., Bochtis, D., 2018. Machine learning in agriculture: a review. *Sensors* 18 (8), 2674.
- Lin, T. Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C. L., 2014. Microsoft coco: Common objects in context. In: *European Conference on Computer Vision*, 740–755.
- Liu, M.Y., Tuzel, O., 2016. Coupled generative adversarial networks. In: *The 29th Conference on Neural Information Processing Systems (NIPS)*, Barcelona, Spain, 1–39.
- Liu, M.Y., Huang, X., Yu, J., Wang, T.C., Mallya, A., 2021. Generative adversarial networks for image and video synthesis: algorithms and applications. *Proc. IEEE* 109 (5), 839–862.
- Liu, B., Tan, C., Li, S., He, J., Wang, H., 2020. A data augmentation method based on generative adversarial networks for grape leaf disease identification. *IEEE Access* 8, 102188–102198.
- Liu, C., Zoph, B., Neumann, M., Shlens, J., Hua, W., Li, L.J., Fei-Fei, L., Yuille, A., Huang, J., Murphy, K., 2018. Progressive neural architecture search. In: *Proceedings of the European Conference on Computer vision (ECCV)*, pp. 19–34.
- Lu, C.Y., Rustia, D.J.A., Lin, T.T., 2019. Generative adversarial network-based image augmentation for insect pest classification enhancement. *IFAC-Papers OnLine* 52 (30), 1–5.
- Lu, Y., Young, S., 2020. A survey of public datasets for computer vision tasks in precision agriculture. *Comput. Electron. Agric.* 178, 105760.
- Lu, Y., Tai, Y.W., Tang, C.K., 2018. Attribute-guided face generation using conditional cycleGAN. In: *Proceedings of the European conference on computer vision (ECCV)*, pp. 282–297.
- Lu, Y., Saeys, W., Kim, M., Peng, Y., Lu, R., 2020. Hyperspectral imaging technology for quality and safety evaluation of horticultural products: A review and celebration of the past 20-year progress. *Postharvest Biol. Technol.* 170, 111318.
- Lucic, M., Kurach, K., Michalski, M., Gelly, S., Bousquet, O., 2019. Are GANs created equal? a large-scale study. In: *The 32nd Conference on Neural Information Processing Systems*, pp. 698–707.
- Luo, Z., Yu, H., Zhang, Y., 2020. Pine cone detection using boundary equilibrium generative adversarial networks and improved yolov3 model. *Sensors* 20 (16), 4430.
- Madsen, S.L., Dyrmann, M., Jørgensen, R.N., Karstoft, H., 2019a. Generating artificial images of plant seedlings using generative adversarial networks. *Biosyst. Eng.* 187, 147–159.
- Madsen, S.L., Mortensen, A.K., Jørgensen, R.N., Karstoft, H., 2019b. Disentangling information in artificial images of plant seedlings using semi-supervised GAN. *Remote Sens.* 11 (22), 2671.
- Maheswari, P., Raja, P., Apolo-Apolo, O.E., Perez-Ruiz, M., 2021. Intelligent fruit yield estimation for orchards using deep learning-based semantic segmentation techniques—a review. *Front. Plant Sci.* 12, 1247.
- Mahlein, A.-K., Kuska, M.T., Behmann, J., Polder, G., Walter, A., 2018. Hyperspectral sensors and imaging technologies in phytopathology: state of the art. *Annu. Rev. Phytopathol.* 56, 535–558.
- Maqsood, M.H., Mumtaz, R., Haq, I.U., Shafi, U., Zaidi, S.M.H., Hafeez, M., 2021. Super resolution generative adversarial network (srgans) for wheat stripe rust classification. *Sensors* 21 (23), 7903.
- Marburg, A., Bigham, K., 2016. Deep learning for benthic fauna identification. In: *OCEANS 2016 MTS/IEEE Monterey*, 1–5.
- Mariani, G., Scheidegger, F., Istrate, R., Bekas, C., Malossi, C., 2018. Bagan: data augmentation with balancing GAN. arXiv preprint arXiv:1803.09655.
- Marino, S., Beausery, P., Smolarz, A., 2020. Unsupervised adversarial deep domain adaptation method for potato defects classification. *Comput. Electron. Agric.* 174, 105501.
- Mavridou, E., Vrochidou, E., Papakostas, G.A., Pachidis, T., Kaburlasos, V.G., 2019. Machine vision systems in precision agriculture for crop farming. *J. Imaging* 5 (12), 89.
- Metz, L., Poole, B., Pfau, D., Sohl-Dickstein, J., 2016. Unrolled generative adversarial networks. arXiv preprint arXiv:1611.02163.
- Minervini, M., Scharr, H., Tsaftaris, S.A., 2015. Image analysis: the new bottleneck in plant phenotyping. *IEEE Signal Process. Mag.* 32 (4), 126–131.
- Minervini, M., Fischbach, A., Scharr, H., Tsaftaris, S.A., 2016. Finely grained annotated datasets for image-based plant phenotyping. *Pattern Recogn. Lett.* 81, 80–89.
- Mirza, M., and Osindero, S. 2014. Conditional generative adversarial nets. arXiv preprint arXiv:1411.1784.
- Miyato, T., Kataoka, T., Koyama, M., Yoshida, Y., 2018. Spectral normalization for generative adversarial networks. arXiv preprint arXiv:1802.05957, 2018.
- Moher, D., Liberati, A., Tetzlaff, J., Altman, D.G., Group, P., 2009. Preferred reporting items for systematic reviews and meta-analyses: the prisma statement. *Ann. Internal Med.* 151 (4), 264–269.
- Monteiro, A., Santos, S., Goncalves, P., 2021. Precision agriculture for crop and livestock farming brief review. *Animals* 11 (8), 2345.
- Naem, M.F., Oh, S. J., Uh, Y., Choi, Y., Yoo, J., 2020. Reliable fidelity and diversity metrics for generative models. In: *Proceedings of the 37th International Conference on Machine Learning* 119, 7176–7185.
- Najafian, K., Ghanbari, A., Stavness, I., Jin, L., Shirdeh, G. H., Maleki, F. A., 2021. Semi-self-supervised learning approach for wheat head detection using extremely small number of labeled samples. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 1342–1351.
- Nasrollahi, K., Moeslund, T.B., 2014. Super-resolution: a comprehensive survey. *Mach. Vis. Appl.* 25 (6), 1423–1468.
- Nazki, H., Lee, J., Yoon, S., Park, D.S., 2019. Image-to-image translation with GAN for synthetic data augmentation in plant disease datasets. *Korean Instit. Smart Media* 8 (2), 46–57.
- Nazki, H., Yoon, S., Fuentes, A., Park, D.S., 2020. Unsupervised image translation using adversarial networks for improved plant disease recognition. *Comput. Electron. Agric.* 168, 105117.
- Nerkar, B., Talbar, S., 2021. Cross-dataset learning for performance improvement of leaf disease detection using reinforced generative adversarial networks. *Int. J. Inform. Technol.* 13 (6), 2305–2312.
- O'Mahony, N., Campbell, S., Carvalho, A., Harapanahalli, S., Hernandez, G. V., Krpalkova, L., Riordan, D., Walsh, J., 2019. Deep learning vs. traditional computer vision. In: *Science and Information Conference*, 128–144.
- Odena, A., Olah, C., Shlens, J., 2017. Conditional image synthesis with auxiliary classifier GANs. In: *International conference on machine learning*, 2642–2651.
- Olatunji, J., Redding, G., Rowe, C., East, A., 2020. Reconstruction of kiwifruit fruit geometry using a cgan trained on a synthetic dataset. *Comput. Electron. Agric.* 177, 105699.
- Oliveira, D.A.B., Pereira, L.G.R., Bresolin, T., Ferreira, R.E.P., Dorea, J.R.R., 2021. A review of deep learning algorithms for computer vision systems in livestock. *Livestock Sci.* 253, 104700.

- Oliveira, A., Pereira, S., Silva, C.A., 2017. Augmenting data when training a cnn for retinal vessel segmentation: How to warp?. In: 2017 IEEE 5th Portuguese Meeting on Bioengineering (ENBENG), pp. 1–4.
- Park, T., Liu, M.Y., Wang, T.C., Zhu, J.Y., 2019. Semantic image synthesis with spatially adaptive normalization. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 2337–2346.
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Kopf, A., Yang, E.Z., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., Chintala, S., 2019. Pytorch: An imperative style, high-performance deep learning Library. In: The 33rd International Conference on Neural Information Processing Systems 721, 8026–8037.
- Paullada, A., Raji, I.D., Bender, E.M., Denton, E., Hanna, A., 2021. Data and its (dis) contents: A survey of dataset development and use in machine learning research. *Patterns* 2 (11), 100336.
- Perez, L., Wang, J., 2017. The effectiveness of data augmentation in image classification using deep learning. arXiv preprint arXiv:1712.04621.
- Perez-García, F., Sparks, R., Ourselin, S., 2021. Torchio: a python library for efficient loading, preprocessing, augmentation and patch-based sampling of medical images in deep learning. *Comput. Methods Programs Biomed.* 208, 106236.
- Qin, L., Zhang, Z., Zhao, H., 2016. Stacking gated neural architecture for implicit discourse relation classification. In: Proceedings of the Conference on Empirical Methods in Natural Language Processing, pp. 2263–2270.
- Quinero-Candela, J., Sugiyama, M., Schwaighofer, A., Lawrence, N.D., 2008. Dataset shift in machine learning. MIT Press.
- Radford, A., Metz, L., Chintala, S., 2015. Unsupervised representation learning with deep convolutional generative adversarial networks. arXiv preprint arXiv:1511.06434.
- Rahman, M.S., Rivera, E., Khomh, F., Gueheneuc, Y.G., Lehnert, B., 2019. Machine learning software engineering in practice: An industrial case study. arXiv preprint arXiv:1906.07154.
- Redmon, J., Farhadi, A., 2018. Yolov3: An incremental improvement. arXiv preprint arXiv:1804.02767.
- Reed, S., Akata, Z., Yan, X., Logeswaran, L., Schiele, B., Lee, H., 2016. Generative adversarial text to image synthesis. In: International conference on machine learning, 1060–1069.
- Ronneberger, O., Fischer, P., Brox, T., 2015. U-net: Convolutional networks for biomedical image segmentation. In: International Conference on Medical image computing and computer-assisted intervention, 234–241.
- Rusu, A.A., Rabinowitz, N.C., Desjardins, G., Soyer, H., Kirkpatrick, J., Kavukcuoglu, K., Pascanu, R., Hadsell, R., 2016. Progressive neural networks. arXiv preprint arXiv:1606.04671.
- Sa, I., Chen, Z., Popovic, M., Khanna, R., Liebisch, F., Nieto, J., Siegart, R., weednet: Dense semantic weed classification using multispectral images and mav for smart farming. *IEEE Robot. Automation Lett.* 3(1), 588–595.
- Salimans, T., Goodfellow, I., Zaremba, V., Cheung, V., Radford, A., Chen, X., 2016. Improved techniques for training GANs. In: 30th Conference on Neural Information Processing Systems.
- Shanmugam, K., Gadhamshetty, V., Tysklind, M., Bhattacharyya, D., Upadhyayula, V.K., 2022. A sustainable performance assessment framework for circular management of municipal wastewater treatment plants. *J. Cleaner Prod.* 339, 130657.
- Shapiro, L.G., Stockman, G.C., 2001. Computer vision. Pearson.
- Shelhamer, E., Long, J., Darrell, T., 2016. Fully convolutional networks for semantic segmentation. *IEEE Trans. Pattern Anal. Mach. Intell.* 39 (4), 640–651.
- Shete, S., Srinivasan, S., Gonsalves, T.A., 2020. Tasselgan: An application of the generative adversarial model for creating field-based maize tassel data. *Plant Phenomics* 2020, 8309605.
- Shorten, C., Khoshgoftar, T.M., 2019. A survey on image data augmentation for deep learning. *J. Big Data* 6 (1), 1–48.
- Siegford, J.M., Berezowski, J., Biswas, S.K., Daigle, C.L., Gebhardt-Henrich, S.G., Hernandez, C.E., Thurner, S., Toscano, M.J., 2016. Assessing activity and location of individual laying hens in large groups using modern technology. *Animals* 6 (2), 10.
- Simard, P.Y., Steinkraus, D., Platt, J.C., 2003. Best practices for convolutional neural networks applied to visual document analysis. Proceedings Seventh International Conference on Document Analysis and Recognition.
- Simonyan, K., Zisserman, A., 2014. Very deep convolutional networks for large-scale image recognition. arXiv preprint arXiv:1409.1556.
- Singh, P., Devi, K.J., Varish, N., 2021. Muzzle pattern-based cattle identification using generative adversarial networks. In: *Soft Computing for Problem Solving*, 13–23.
- Singh, A.K., Ganapathy Subramanian, B., Sarkar, S., Singh, A., 2018. Deep learning for plant stress phenotyping: trends and future perspectives. *Trends Plant Sci.* 23 (10), 883–898.
- Spurr, A., Aksan, E., Hilliges, O., 2017. Guiding infogan with semi supervision. In: Joint European Conference on Machine Learning and Knowledge Discovery in Databases, pp. 119–134.
- Srivastava, A., Valkov, L., Russell, C., Gutmann, M.U., Sutton, C., 2017. Veegan: Reducing mode collapse in gans using implicit variational learning. *Advances in Neural Information Processing Systems*.
- Sun, C., Shrivastava, A., Singh, S., Gupta, A., 2017. Revisiting unreasonable effectiveness of data in deep learning era. In: Proceedings of the IEEE International Conference on Computer Vision, pp. 843–852.
- Sun, R., Zhang, M., Yang, K., Liu, J., 2020. Data enhancement for plant disease classification using generated lesions. *Appl. Sci.* 10 (2), 466.
- Taylor, L., Nitschke, G., 2018. Improving deep learning with generic data augmentation. In: *IEEE Symposium Series on Computational Intelligence (SSCI)*, pp. 1542–1547.
- Thomas, S., Kuska, M.T., Bohnenkamp, D., Brugger, A., Alisaac, E., Wahabzada, M., Behmann, J., Mahlein, A.K., 2018. Benefits of hyperspectral imaging for plant disease detection and plant protection: a technical perspective. *J. Plant Dis. Prot.* 125 (1), 5–20.
- Tian, Y., Yang, G., Wang, Z., Li, E., Liang, Z., 2019. Detection of apple lesions in orchards BASED on deep learning methods of cyclegan and yolov3-dense. *J. Sens.* 2019, 7630926.
- Tomczak, J.M., 2022. Why Deep Generative Modeling? pp. 1–12. Springer International Publishing, Cham, 2022.
- Tran, N.T., Tran, V.H., Nguyen, N.B., Nguyen, T.K., Cheung, N.M., 2021. On data augmentation for GAN training. *IEEE Trans. Image Process.* 30, 1882–1897.
- Tzeng, E., Hoffman, J., Saenko, K., Darrell, T., 2017. Adversarial discriminative domain adaptation. In: Proceedings of the IEEE conference on computer vision and pattern recognition, pp. 7167–7176.
- Vandenhende, S., Georgoulis, S., Van Gansbeke, W., Proesmans, M., Dai, D., Van Gool, L., 2021. Multitask learning for dense prediction tasks: A survey. *IEEE Trans. Pattern Anal. Mach. Intell.* <https://doi.org/10.1109/TPAMI.2021.3054719>.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I., 2017. Attention is all you need. 31st Conference on Neural Information Processing Systems.
- Vougioukas, S.G., 2019. Agricultural robotics. *Ann. Rev. Control, Robot., Autonomous Syst.* 2, 365–392.
- Wang, Z., Bovik, A.C., Sheikh, H.R., Simoncelli, E.P., 2004. Image quality assessment: from error visibility to structural similarity. *IEEE Trans. Image Process.* 13 (4), 600–612.
- Wang, Z., Chen, J., Hoi, S.C., 2020. Deep learning for image super-resolution: a survey. *IEEE Trans. Pattern Anal. Mach. Intell.* 43 (10), 3365–3387.
- Wang, X., Girshick, R., Gupta, A., He, K., 2018a. Non-local neural networks. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 7794–7803.
- Wang, C., Xiao, Z., 2021. Lychee surface defect detection based on deep convolutional neural networks with gan-based data augmentation. *Agronomy* 11 (8), 1500.
- Wang, M., Yang, G.Y., Li, R., Liang, R.Z., Zhang, S.H., Hall, P.M., Hu, S.M., 2019. Example-guided style-consistent image synthesis from semantic labeling. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 1495–1504.
- Wang, X., Yu, K., Wu, S., Gu, J., Liu, Y., Dong, C., Qiao, Y., Change Loy, C., 2018b. Esrgan: Enhanced super-resolution generative adversarial networks. Proceedings of the European Conference on Computer Vision (ECCV) Workshops.
- Wang, X., Xie, L., Dong, C., Shan, Y., 2021. Real-esrgan: Training real-world blind super-resolution with pure synthetic data. In: Proceedings of the IEEE/CVF International Conference on Computer Vision, pp. 1905–1914.
- Wen, J., Shi, Y., Zhou, X., Xue, Y., 2020. Crop disease classification on inadequate low-resolution target images. *Sensors* 20 (16), 4601.
- Westwood, J.H., Charudattan, R., Duke, S.O., Fennimore, S.A., Marrone, P., Slaughter, D. C., Swanton, C., Zollinger, R., 2018. Weed management in 2050: Perspectives on the future of weed science. *Weed Sci.* 66 (3), 275–285.
- Wiatrak, M., Albrecht, S. V., Nystrom, A., 2019. Stabilizing generative adversarial networks: a survey. arXiv preprint arXiv:1910.00927.
- Wong, S.C., Gatt, A., Stamatescu, V., McDonnell, M.D., 2016. Understanding data augmentation for classification: when to warp? International Conference on Digital Image Computing: Techniques and Applications (DICTA).
- Wu, Q., Chen, Y., Meng, J., 2020. Dcgan-based data augmentation for tomato leaf disease identification. *IEEE Access* 8, 98716–98728.
- Wu, Y., Kirillov, A., Massa, F., Lo, W.Y., Girshick, R., 2019. Detectron2. <https://github.com/facebookresearch/detectron2>.
- Wurtz, K., Camerlink, I., DEath, R. B. A. P., Fernandez, T., Norton, J., Steibel, Siegford, J., 2019. Recording behaviour of indoor-housed farm animals automatically using machine vision technology: systematic review. *PLoS ONE* 14(12), e0226669.
- Xu, M., Yoon, S., Fuentes, A., Yang, J., Park, D. S., 2022. Style-consistent image translation: a novel data augmentation paradigm to improve plant disease recognition. *Front. Plant Sci.* 12, 773142-773142.
- Xu, Q., Huang, G., Yuan, Y., Guo, C., Sun, Y., Wu, F., Weinberger, K., 2018. An empirical study on evaluation metrics of generative adversarial networks. 6th International Conference on Learning Representations.
- Xu, L., Skoularidou, M., Cuesta-Infante, A., Veeramachaneni, K., 2019. Modeling tabular data using conditional gan. 33rd Conference on Neural Information Processing Systems.
- Yang, X., Guo, M., Lyu, Q., Ma, M., 2021. Detection and classification of damaged wheat kernels based on progressive neural architecture search. *Biosyst. Eng.* 208, 176–185.
- Yi, X., Walia, E., Babyn, P., 2019. Generative adversarial network in medical imaging: A review. *Med. Image Anal.* 58, 101552.
- Yi, Z., Zhang, H., Tan, P., Gong, M., 2017. Dualgan: Unsupervised dual learning for image-to-image translation. In: Proceedings of the IEEE international conference on computer vision, pp. 2849–2857.
- Yoo, J., Ahn, N., Sohn, K.A., 2020. Rethinking data augmentation for image super-resolution: A comprehensive analysis and a new strategy. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 8375–8384.
- Yuwana, R.S., Fauziah, F., Heryana, A., Krisnandi, D., Kusumo, R.B.S., Pardede, H.F., 2020. Data augmentation using adversarial networks for tea disease detection. *Jurnal Elektronika dan Telekomunikasi* 20 (1), 29–35.
- Zeng, M., Gao, H., Wan, L., 2021. Few-shot grape leaf diseases classification based on generative adversarial network. In: *Journal of Physics Conference Series* 1883, 012093.
- Zeng, Q., Ma, X., Cheng, B., Zhou, E., Pang, W., 2020. Gans-based data augmentation for citrus disease severity detection using deep learning. *IEEE Access* 8, 172882–172891.

- Zhai, X., Kolesnikov, A., Houlsby, N., Beyer, L., 2021. Scaling vision transformers. arXiv preprint arXiv: 2106.04560. <https://arxiv.org/abs/2106.04560>.
- Zhang, Y., Chen, Z., Gao, X., Song, W., Xiong, Q., Hu, J., Zhang, Z., 2021a. Plant disease detection using generated leaves based on DoubleGAN. *IEEE/ACM Trans. Comput. Biol. Bioinf.* 19 (3), 1817–1826.
- Zhang, H., Goodfellow, I., Metaxas, D., Odena, A., 2019a. Self-attention generative adversarial networks. In: *The 36th International Conference on Machine Learning* 97, 7354–7363.
- Zhang, M., Liu, S., Yang, F., Liu, J., 2019b. Classification of canker on small datasets using improved deep convolutional generative adversarial networks. *IEEE Access* 7, 49680–49690.
- Zhang, Q., Liu, Y., Gong, C., Chen, Y., Yu, H., 2020. Applications of deep learning for dense scenes analysis in agriculture: A review. *Sensors* 20 (5), 1520.
- Zhang, J., Rao, Y., Man, Y., Jiang, Z., Li, S., 2021b. Identification of cucumber leaf diseases using deep learning and small sample size for agricultural internet of things. *Int. J. Distrib. Sens. Netw.* 17 (4), 2–13.
- Zhang, H., Xu, T., Li, H., Zhang, S., Wang, X., Huang, X., Metaxas, D.N., 2017. Stackgan: text to photo-realistic image synthesis with stacked generative adversarial networks. In: *Proceedings of the IEEE international conference on computer vision*, pp. 5907–5915.
- Zhang, Y., Yang, Q., 2021. A survey on multi-task learning. *IEEE Trans. Knowl. Data Eng.* <https://doi.org/10.1109/TKDE.2021.3070203>.
- Zhang, G., Yang, L., Sun, C., Zhou, X., Zhou, Q., Li, M. Bi, Guo, J., 2021c. Shrimp egg counting with fully convolutional regression network and generative adversarial network. *Aquacult. Eng.* 94, 102175.
- Zhao, Y., Chen, Z., Gao, X., Song, W., Xiong, Q., Hu, J., Zhang, Z., 2021b. Plant disease detection using generated leaves based on doublegan. *IEEE/ACM Trans. Comput. Biol. Bioinf.* <https://doi.org/10.1109/TCBB.2021.3056683>.
- Zhao, J., Li, Y., Zhang, F., Zhu, S., Liu, Y., Lu, H., Ye, Z., 2018. Semi-supervised learning-based live fish identification in aquaculture using modified deep convolutional generative adversarial networks. *Trans. ASABE* 61 (2), 699–710.
- Zhao, L., Zhang, Z., Chen, T., Metaxas, D., Zhang, H., 2021a. Improved transformer for high-resolution gans. *35th Conference on Neural Information Processing Systems*.
- Zhou, S., Gordon, M., Krishna, R., Narcomey, A., Fei-Fei, L. F., Bernstein, M., 2019. Hype: A benchmark for human eye perceptual evaluation of generative models. In: *The 33rd International Conference on Neural Information Processing Systems* 310, 3449–3461.
- Zhu, Y., Aoun, M., Krijn, M. Vanschoren, J., Campus, H.T., 2018. Data augmentation using conditional generative adversarial networks for leaf counting in arabidopsis plants. In: *British Machine Vision Conference, Workshop on Computer Vision Problems in Plant Phenotyping*.
- Zhu, F., He, M., Zheng, Z., 2020. Data augmentation using improved cdcgan for plant vigor rating. *Comput. Electron. Agric.* 175, 105603.
- Zhu, J.Y., Park, T., Isola, P., Efros, A.A., 2017. Unpaired image-to-image translation using cycle-consistent adversarial networks. In: *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pp. 2223–2232.