

# BugNIST – a Large Volumetric Dataset for Object Detection under Domain Shift

Patrick Møller Jensen<sup>✉</sup>, Vedrana Andersen Dahl<sup>✉</sup>, Rebecca Engberg<sup>✉</sup>,  
Carsten Gundlach<sup>✉</sup>, Hans Marin Kjer<sup>✉</sup>, and Anders Bjorholm Dahl<sup>✉</sup>

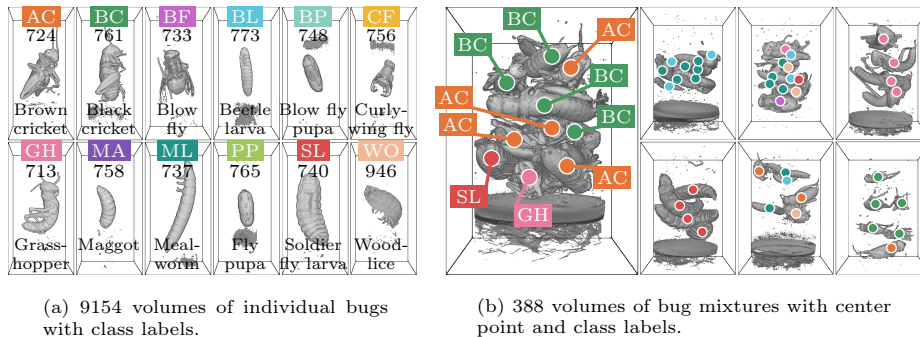
Technical University of Denmark, Kgs. Lyngby, Denmark  
{patmjen,vand,reen}@dtu.dk, cagu@fysik.dtu.dk, {hmjk,abda}@dtu.dk

**Abstract.** Domain shift significantly influences the performance of deep learning algorithms, particularly for object detection within volumetric 3D images. Annotated training data is essential for deep learning-based object detection. However, annotating densely packed objects is time-consuming and costly. Instead, we suggest training models on individually scanned objects, causing a domain shift between training and detection data. To address this challenge, we introduce the BugNIST dataset, comprising 9154 micro-CT volumes of 12 bug types and 388 volumes of tightly packed bug mixtures. This dataset is characterized by having objects with the same appearance in the source and target domains, which is uncommon for other benchmark datasets for domain shift. During training, individual bug volumes labeled by class are utilized, while testing employs mixtures with center point annotations and bug type labels. Together with the dataset, we provide a baseline detection analysis, with the aim of advancing the field of 3D object detection methods.

**Keywords:** Volumetric Dataset, Benchmark, Volumetric Object Detection, Domain Shift.

## 1 Introduction

Our work on domain shift in volumetric 3D images is motivated by the need for labeled data to train supervised deep learning models for volumetric imaging. The problem that we aim to solve is object detection and classification. We propose to label images of objects scanned as isolated entities as the source domain and use these as a basis to train models for object detection and classification in a complex context of mixed objects and other materials as the target domain. The effort needed for obtaining labels in the two domains is significantly different. If objects are isolated, they can be automatically labeled whereas mixed objects require expensive manual labeling. Automating the labeling based on isolated objects, however, leads to a domain shift between the data in the source domain for training and the data in the target domain for detection and classification. This domain shift is special because the appearance of the objects is the same



**Fig. 1:** Overview of the two parts of the dataset. In (a), representative individual bug volumes are displayed with a class abbreviation, name, and number of volumes. In (b), a few volumes of bug mixtures are shown. Each bug is annotated with bug class and center point.

in the two domains, but the surrounding context is different. We use the term *context shift* for this type of domain shift.

To promote the development of new deep learning-based methods that can handle context shift, we contribute the volumetric dataset called BugNIST<sup>1</sup>. The dataset contains 9154 micro-CT volumes with individual bugs and 388 micro-CT volumes containing a mixture of several bugs as well as other materials such as leaves, wood shavings, *etc.*; see Fig. 1. The challenge is to use the individually scanned bugs to detect and classify the bugs in the mixtures. The dataset is intended to be a relatively simple test case for starting to develop methods for detecting objects in volumetric images under context shift, and bugs are chosen due to their complex shape and *not* for solving a problem in entomology.

As a basis for training an accurate deep learning model, a training dataset for object detection and classification must capture the appearance distribution of the object classes that are in the dataset. Since the appearance of a CT-scanned object changes minimally if it is scanned individually or as part of a mixture, the appearance distribution can be covered by only considering the individual scans. This is different from 2D photos. Here additional variations in scale, viewpoint, and illumination influence the appearance of objects, and it would be more difficult to acquire representative images of individual objects for a dataset based on 2D photos.

Since BugNIST is a test case for context shift, we wanted to ensure that each class is covered well. By scanning systematically arranged bugs, we have generated 9154 volumes of individual bugs, each labeled with its corresponding class. Furthermore, we scanned 388 mixed volumes, each containing between 4 and 45 bugs alongside other materials. For evaluation purposes, each bug in the mixtures is manually annotated with its center position and class label. In total,

<sup>1</sup> <https://abdahl.github.io/bugnist/>

9542 volumes make up the BugNIST dataset. The numbers in each class are shown in Fig. 1. Since the bugs in each class are the same age, we expect to cover the biological variation sufficiently to be representative of the class.

Solutions to the type of context shift, that we present here, namely a domain shift where only the context changes, and not the appearance of the objects, have wide applications in volumetric imaging. The field of volumetric imaging is dominated by medical applications [39], but there is a growing use beyond the medical area based on *e.g.* micro-CT [58, 63]. In all these application areas, there is a need to employ supervised deep learning to solve problems that require quantification of depicted structures. When the use case is novel and there is no directly applicable training data available — a typical scenario in volumetric imaging of *e.g.* materials microstructure, but also in specialized biomedical applications — it is necessary to *create* training data. If labor-intensive manual labeling can be limited or avoided, there is potentially a huge gain in productivity. Examples of applications are detecting and classifying objects in scans of marine sediments such as microplastics [55, 68], museum specimens [20], or detecting threat items in baggage security screening [56].

In volumetric scans, manual labeling is significantly more demanding than in 2D. For instance, annotating a bounding box of a single object in 3D requires traversing slices of the volume and making a mental model of the object. If the volume contains densely packed complex objects, determining their bounding boxes becomes difficult. Further, bounding boxes may overlap to such an extent that they stop being a useful way to locate objects. This is why we employ center points as the annotation format in the mixtures.

The name BugNIST is a homage to the influential MNIST dataset [29] which has played a pivotal role in the development and evaluation of deep learning methods. Similar to the way that MNIST has been used, BugNIST is intended to be a benchmark dataset that represents a well-defined problem, where the development of solutions will make a basis for solutions that can advance the field. Specifically for BugNIST, solutions will enable object detection and classification with significantly less effort put into creating the data foundation for training. This calls for methods that can learn to predict objects in a complex context from volumes of individual objects, namely solving the *context shift* problem. To our knowledge, no existing datasets have this property.

The term *context shift* refers to a shift in the data distribution between training and inference. Typical shifts in data distributions arise from changes in data acquisition systems, *e.g.* the use of different scanners, or from variations in the objects of interest between datasets. Such problems have been addressed using domain adaptation [18, 59]. In our setting, there is almost no change in the appearance of the bugs between volumes of individual bugs and mixtures, but the change is in what surrounds the bugs, *i.e.* the change is in the context.

We hypothesize that methods capable of learning under *context shift* might also prove capable of learning from very few examples, similar to how humans learn [67]. Instead of relying on seeing objects in every imaginable context, such methods could instead isolate objects from their contexts. If this hypothesis holds

true, learning under context shift may reduce the need for extensive training data and the associated training costs. This is in contrast to the current trend, where larger models and datasets as seen in LLMs [69] and computer vision [28, 35], lead to better performance, but also limit who can effectively train such models.

*Contributions* We contribute BugNIST, a new volumetric image dataset for benchmarking domain shift. In this paper, we use BugNIST for benchmarking detection methods under context shift, trained only on individual bugs and tested on mixtures. We provide several baselines and investigate different training strategies for existing volumetric detection methods.

Our results show that volumetric object detection under context shift is difficult. We obtain reasonable results in localizing the bugs, but when a class label is assigned, there is a significant drop in performance. This is even though the classification of individual bugs is easy. Our experiments demonstrate great potential for developing new approaches for 3D object detection of complex objects.

## 2 Related work

In computer vision, the term 3D data refers to different data types, including point clouds, surfaces, volumes, and depth images. Here, we are only concerned with 3D volumes, *i.e.* image data represented on a voxel grid.

BugNIST is created to advance deep learning methods for object detection and classification under domain shift. Especially in medical imaging, analysis of 3D volumes is in great demand. A review of medical challenges illustrates this [36]. Their study shows that detection makes up 7% of the tasks, classification makes up 10% of the tasks, whereas segmentation accounts for 70% of the tasks. Detection is important for example for diagnosis [7, 19, 27, 49, 71]. Outside the medical domain, object detection problems include baggage security screening [15, 61, 62], although the problem and the available data are still primarily in the 2D domain [56], food inspection [12] such as fruit quality [24], grain classification [6], foreign object detection [13], biological applications such as analysis of weeds [64] and other uses driven by developments in X-ray micro-CT imaging [42], *etc.*

The domain shift problem presented in BugNIST, where only the context changes, and not the appearance of the objects of interest, is not widely studied. Most research on domain shifts focuses primarily on problems where the appearance of objects changes. Solutions focus on domain adaptation, where labeled data in the source domain are available together with data without labels from the target domain [40], or domain generalization, where the inference domain is not known at training time [37, 60]. We focus on domain adaptation, which is also the typical use case for volumetric imaging.

Several widely used datasets are available for studying and developing methods for domain shift. For 2D image classification problems the datasets PACS [30], Digits-DG [70], Office-Home [57], and DomainNet [41] developed to classify images from domains such as photos, art paintings, sketches, *etc.* Domain shift between artificially generated and real images *e.g.* for segmentation is another area that receives much attention [48] with standard datasets such as

SYNTHEIA [47], CITYSCAPES [9], and GTA-5 [45], or completely synthetic domain shift like the SHIFT dataset [52] for autonomous driving.

In 3D, domain shift is often studied in medical applications where the same type of object is imaged with different scanning modalities [32]. It can be MRI (magnetic resonance imaging), CT, OCT (optical coherence tomography), *etc.* Standard datasets include the CrossMoDA 2021 dataset [11] and the BRATS dataset [3], both of different MRI modalities. There are many other datasets used for domain adaptation for volumetric medical imaging [32], but all for adapting between modalities or between the same modality but different scanners [18].

The structure of the imaged objects in BugNIST is different from most existing volumetric deep learning benchmark datasets. The morphological complexity of bugs is, however, representative of volumetric imaging data of *e.g.* 3D microstructure, where there is much less use of deep learning than in *e.g.* the medical domain [63].

Platforms like Zenodo [14], TomoBank [10], or the Human Organ Atlas [58] also contain a large amount of volumetric image datasets. Different from *e.g.* medical challenges, this data generally comes with no clear problem formulation, well-defined tasks or even labels, that allow for a comparison of methods. Further, their large volume sizes are difficult to access. It can also be difficult to judge the complexity of the data and if deep learning is needed. Some images can easily be analyzed using a threshold whereas others are complex and can greatly benefit from deep learning. Therefore, most of this data is not attractive to use as a benchmark for developing new deep learning methods. Some datasets are emerging, *e.g.* the vessel segmentation challenge [23], but they remain exceptions.

There is extensive research in deep learning-based object detection, but primarily in 2D. Most methods either use a two-stage detector such as R-CNN [16, 17, 44], with a region proposal and a classifier to choose among the proposed regions, or a one-stage detector such as YOLO [43], where regions and classes are proposed in a single network.

Object detection in volumetric images is less common. One approach has been to apply a 3D U-Net [8] segmentation network and use each connected component of the foreground label class as a detection [22]. An example of using this approach is lung nodule detection in [65]. Employing segmentation for object detection comes with the problem of separating objects that may be connected or combining separated segments into one object. To avoid this problem, Jaeger et al. [22] propose the Retina U-Net: a one-shot 3D detection algorithm that combines the Retina Net [33] and U-Net [46] giving detections with 3D bounding boxes. Despite the good performance of Retina U-Net, changing to a new 3D detection problem requires a time-consuming method configuration. This has been addressed in nnDetection [4]. Inspired by the nnU-Net (no new U-Net) [21], they propose a no new 3D detection method. On the BugNIST data, we test detection with 3D U-Net [8], the nnDetection method [4], and a 3D adaption of the Faster R-CNN model [44].

Benchmark datasets for volumetric object detection are primarily medical. Examples include Baumgartner et al. [4], which uses 13 datasets to develop the

nnDetection method. Two of these datasets were originally targeted 3D detection and include the LUNA16 for nodule detection [50], which is a subset of the LIDC dataset [2] and contains 888 CT scans of lungs, and the ADAM dataset [53] for detecting aneurysms in MRI scans. ADAM is made up of 254 MRI scans. The rest of the detection datasets are medical segmentation datasets such as data from the Medical Segmentation Decathlon [51], the rib fracture detection [26], the PROSTATEx dataset [1], and others. The segmentation datasets have been transformed into detection data by identifying center positions and bounding boxes for the segments. Within baggage security screening, some 3D datasets are obtainable, for example, the ALERT datasets used in [62]. Due to the sensitive nature of the problem, these datasets have copyright restrictions and limited availability under a non-disclosure agreement.


### 3 Data acquisition and processing

Fig. 1 gives an overview of the BugNIST dataset. Each individual bug volume contains exactly one bug located near the center of the volume. The total of 9154 individual bug volumes have between 713 and 946 volumes of each class distributed across 12 bug classes (species and stages). Bug mixtures contain several bugs, leaves, cotton, shredded paper, and wood shaving. There are in total 388 mixtures. The mixtures contain 4 to 45 bugs, totaling 5011 bugs with an average of 13 bugs per mixture, see Fig. 2. In bug mixtures, we have manually marked the center positions of each bug and provided it with a class label.

All scans were acquired using a laboratory micro-CT scanner (Nikon metrology scanner) that was set to record volumes of 2000 voxels cubed with an isotropic voxel size of 26.62  $\mu\text{m}$ . Being able to acquire such large volumes allowed us to scan multiple bugs and bug mixtures at once, and then crop them into smaller volumes. The dataset was collected through nearly 250 scans, each taking 30 minutes, over 90 days. Care has been taken that the scanning parameters stay unchanged for all scans, and we have calibrated the scanner daily.

All bugs were placed in tubes, either as individuals or in mixtures. We used cotton, which is almost transparent to X-rays, to separate the bugs, either as individual bugs or mixtures. For the mixtures, we varied the difficulty of the bug detection task by adding leaves, shredded paper, or wood shavings to some of the mixtures, and also varied the number of bugs, their class, and the packing density.

The cylindrical shape of the tubes was easy to segment using a circular Hough-transform [66] and sparse layered graphs [25], leaving us with a cylinder of only bugs and cotton. Due to the low intensity of cotton compared to the bugs, we could easily divide each cylinder into volumes containing one bug each. The volumes have been manually inspected and cleaned to correct cases where something went wrong, *e.g.* parts of another bug showing in the periphery of the volume. The final volumes of scans of individual bugs are  $450 \times 450 \times 900$  voxels, the size which was chosen to accommodate the biggest bugs. The mixtures were segmented similarly and the resulting volumes are  $650 \times 650 \times 900$  voxels. All



**Fig. 2:** Summary of volumes from scans of bug mixtures.

volumes are stored as 8-bit unsigned integers with their original spatial resolution of  $26.62 \mu\text{m}$  (physical voxel size). Our dataset comprises 1.8 trillion voxels, approximately a quarter of the RGB values in ImageNet.

The bugs scanned for the BugNIST dataset are bred as pet food and have been frozen before being packed for scanning. Despite careful packing, some scanned bugs may appear damaged, *e.g.* a few crickets have broken off antennas or miss a leg. We have chosen to accept these defects because they reflect some variation in a population and add to the challenge of the dataset.

To create additional lightweight datasets, we created four downsized versions of all volumes. The individual bug volumes, originally 900 voxels long, have lengths of 512, 256, 128, and 64 voxels in the downscaled versions, and all the other lengths are scaled accordingly. The voxel size of the coarsest data is  $374.3 \mu\text{m}$ , which is a scaling factor of approximately 14 compared to the original scan.

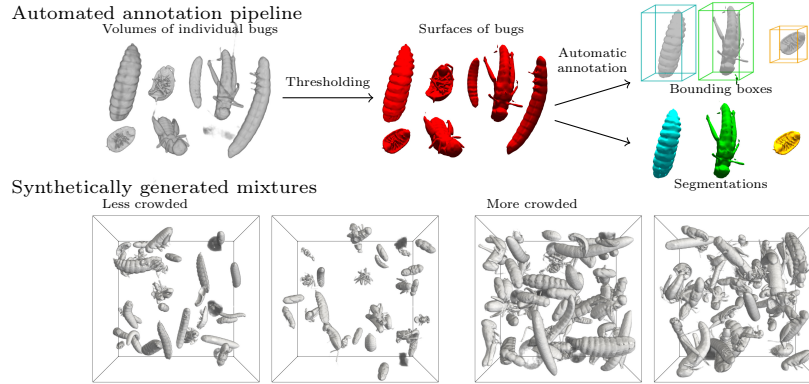
## 4 Experiments

To establish baselines for volumetric detection under context shift (domain shift with no change in object appearance), we have selected and adapted three existing detection models. Furthermore, we have applied three training strategies, yielding a total of nine baseline models. The diversity in model choices and training strategies contributes to a comprehensive set of baseline models that provide a strong foundation for object detection under context shift. We briefly describe the selected methods and how they were trained. Full details are provided in the supplementary material. To determine the difficulty of differentiating between the individual bugs, we also performed a classification experiment which is described in the supplementary material.

All experiments are carried out using the x128 data, which is a compromise between memory needs and the preservation of structural details of the bugs.

### 4.1 Training strategies for detection under context shift

To provide a simple baseline investigation for the BugNIST dataset, we investigate some strategies for solving the detection problem under context shift. Since the purpose of our paper is to initiate research in domain adaptation methods



**Fig. 3:** The top line shows the automated annotation pipeline for individual bugs. By thresholding, bounding boxes and segmentation masks are easily obtained. The bottom shows the synthetic mixtures used for the more advanced training strategy.

for volumetric data under context shift, this baseline study is intended as an offset for future developments based on the BugNIST dataset. Our first attempt is to train models from the volumes of individual bugs. Our two other attempts are to use a data augmentation strategy by creating synthetic mixtures. Data augmentation has previously been shown to greatly benefit domain shift problems [60]. The automated annotation of the BugNIST data and the data that we use for training detection algorithms are illustrated in Fig. 3.

*Single bugs* As a first stage, we train the models to detect the bugs in the individual scans. For models requiring ground truth segmentation labels, those are generated by thresholding the volumes and only keeping the largest connected component. The value of the class label then replaces the foreground voxels. This creates segmentation labels where each voxel has one of 13 values including background or one of the twelve bugs. For models requiring ground truth bounding boxes, those are extracted from each connected component. This process of preparing the training data using this simplest strategy is illustrated in the top line of Fig. 3.

*Synthetic mixtures* As a second stage, we train the models on artificial mixes. These are generated by randomly sampling individual bug scans and adding them to a common volume. We start with an empty volume of  $256^3$  and we set a random target number of bugs drawn from a normal distribution,  $\mathcal{N}(32, 6)$ . A random bug from a random class (both uniform distributions) is drawn. To create variation, the selected sample is flipped around a random axis, rotated a random multiple (0-3) of 90 degrees around the first axis, and then transposed randomly. The bug volume is intensity normalized, smoothed with a Gaussian filter with  $\sigma = 0.5$ , and thresholded at 0.2. We draw a random voxel position



from the volume and only insert the bug if its bounding box does not exceed the volume boundaries nor overlaps with existing labeled voxels. We attempt a maximum of 1000 random positions to achieve a valid insertion before the bug is discarded. In total 300 synthetic mixtures are created.

*Crowded synthetic mixtures* Finally, since the bugs are packed very close together in the real scans, we create crowded synthetic mixtures. Here the bugs are packed closer by reducing the volume size to  $200 \times 200 \times 200$ , adjusting the target number of bugs to  $\mathcal{N}(64, 8)$ , and relaxing the insertion criteria such that no labeled voxels can be overwritten. This effectively allows for a larger overlap between bounding boxes and that bugs can touch each other. Again, a total of 300 synthetic mixtures are created.

## 4.2 Baseline detection models

Despite important uses of detecting objects in 3D volumes, the number of methods addressing volumetric object detection is much smaller compared to 2D, and only a few methods are readily available for volumes. So, to establish baselines for detection under context shift, we have selected and adapted three detection approaches. They vary in how they utilize training data; relying on either ground truth segmentation labels, ground truth bounding boxes, or both.

*Detection U-Net* Inspired by [4] we use U-Net [8] as one baseline detection model. U-Net is a dominating model for volumetric segmentation, making it a natural choice for adapting into a baseline detection model. Detection U-Net is trained on volumes accompanied by ground truth segmentation labels. It is first trained to predict the segmentation of the single bugs. Then, starting from the obtained weights, it is trained to predict the segmentation of the synthetic mixtures. Finally, again starting from the previous weights, it is trained to predict the segmentation in the crowded synthetic mixtures. For all training runs, we augment the images with random rotations, translations, geometric scaling, and intensity scaling. We train the models using the sum of Dice and cross entropy loss. During training we hold out a validation split which we evaluate on after every epoch. At the end of training, the model which scored the highest on the validation split is kept. Like [4], we extract bounding boxes by thresholding the softmax probabilities (at 0.75) and then extracting the connected components for each channel. Only components larger than 125 ( $5^3$ ) voxels are kept. We use the U-Net implementation from MONAI [5].

*Faster R-CNN* As an example of region proposal networks, we have chosen Faster R-CNN [44] with the modifications from [31]. We have adapted the implementation from torchvision [54] to allow for 3D detection. We train in the same sequence as the U-Net: single bugs, synthetic mixtures, and crowded synthetic mixtures. We use the same augmentations as for the U-Net. In addition, we also randomly perturb the bounding box corners. The model is only trained on bounding boxes and does not make use of the voxel labels. The model is

trained with smooth L1 loss for bounding box regression and cross-entropy for objectness and class predictions. For full details, we refer to [31, 44]. As with U-Net, we hold out a validation split which is evaluated after each epoch. After training, the model scoring the highest on the validation split is kept. For inference, only bounding boxes with an objectness score over 0.05 are kept.

*nnDetection* The third detection model, nnDetection, trains the RetinaUNet [22] detection model using a combination of segmentation masks and bounding boxes. The bounding box branch is trained using a generalized IoU loss for regression and cross-entropy for classification. The voxel segmentation branch is trained with a sum of Dice and cross-entropy loss. As nnDetection is a self-configuring framework, we refer to [4, 21] for further training and model details. We simply provide the data as instructed. We train a separate model for single bugs, synthetic mixtures, and crowded synthetic mixtures. nnDetection trains using 5-fold cross-validation, and the final predictions are made using an ensemble over the best model from each fold [4].


### 4.3 Evaluation

We assess model performance by evaluating their detection accuracy on real scans of densely packed bugs, using manually annotated center points. These annotations are matched with predicted bounding box centers based on Euclidean distance, optimized using Munkres’ algorithm [38]. Matches are only accepted if the annotated point falls within a  $1.5\times$  inflated bounding box, as illustrated in Fig. 4. This inflation ensures a fair assessment for detections like in Fig. 4(b). This, combined with matching on the bounding box centers, helps ensure that large boxes are not favored since it places more weight on the location of the boxes. We compute precision to avoid favoring overdetectors, recall to measure the proportion of correct matches, and F1-score to combine the two. Further, we first compute the metrics purely from the point matches without considering classes — thereby evaluating if the models can detect the presence of bugs vs. background material. Next, we recompute the matches where we only allow matches between points of equal class.

## 5 Results

The results are summarized in Tab. 1, example results are shown in Fig. 5, and the distributions of the metrics without class information are shown in Fig. 6. U-Net performs relatively well, independent of what it was trained on, whereas Faster R-CNN performs worse on crowded mixtures and nnDetection on single bugs. If class labels are required, the performance is generally low.

The distributions in Fig. 6 show that U-Net has a narrower spread around the mean for the F1-score, whereas the two others are spread more out. Wrt. the precision, many detections obtained by the nnDetection method are 1, but have a lower recall, meaning that those that it predicts are correct, but there



**Fig. 4:** Matching criterion for bounding box to center point. The circle shows the manually annotated center point, the predicted bounding box is shown with the full line, and the inflated bounding box is shown with the dashed line. In (a), the match is accepted by the predicted bounding box, in (b), the match is also accepted due to the inflated bounding box, whereas in (c), there is no match. Note that this illustration is 2D, but the detection is done in 3D.

are many potential detections that it misses. It is the opposite for the Faster R-CNN, where many datasets get a recall of 1, but a lower precision, meaning that it detects many of the bugs, but also has many false detections.

To investigate the capabilities of the three methods without the domain shift between the synthetic and real mixtures, we included the results in Tab. 2. This table shows the results of the validation data for training the models on the volumes of individual bugs and the two synthetic mixtures with the class labels. Since we have the bounding boxes for these datasets, we compute the mean average precision (mAP), similar to, *e.g.*, COCO detection [34]. We compute mAP with the IoU threshold running from 0.5 to 0.95 (shown as mAP) and fixed at 0.5 (mAP@50). We note that the values are not directly comparable with those in Tab. 1. We show example detections in the supplementary. The U-Net generally has high performance, Faster R-CNN has high performance for single bugs but has difficulty with the mixtures, whereas nnDetection performs better in the mixtures than for the individual bugs. Generally, the detection performance of bugs is far from perfect and can be improved, also when training and testing within the same domain.

## 6 Discussion

The domain shift introduced by shifting the context surrounding the bugs significantly increases the difficulty of the analysis problem. While U-Net and nnDetection achieve high performance on their training tasks (at least for synthetic mixtures), they struggle when transferred to real densely packed scans. It is also interesting that the domain shift has such an effect on the ability of the models to infer the correct class even when accurately localizing the bug. This is surprising since detecting the bugs in the training domain can be done with high accuracy, as shown in Tab. 2, and pure classification (see supplementary) gave an accuracy of 0.97. Despite the models for classification being different than the detection

**Table 1:** Detection results on mixed scans. We show mean and std. for each metric computed over the evaluation images. Higher is better. Without class info refers to not considering the bug label and with class info expects the bug type to be classified as well as being detected.

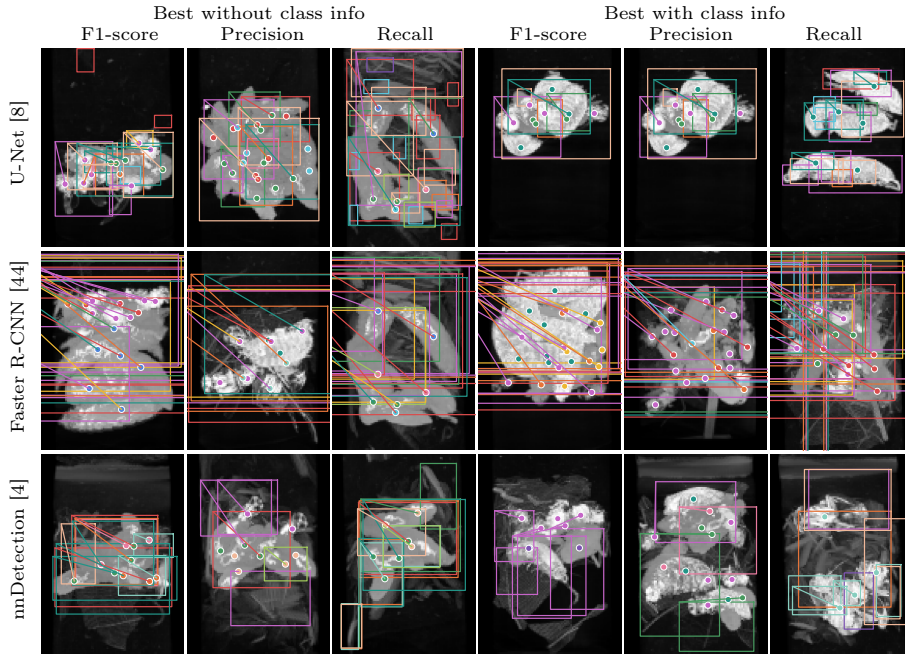
Model	- Trained on	Without class info.			With class info.		
		F1-Score	Precision	Recall	F1-Score	Precision	Recall
U-Net [8]	- Single bugs	0.59±0.18	0.68±0.17	0.60±0.25	0.11±0.16	0.13±0.19	0.11±0.16
	- Synth. mixes	0.60±0.12	0.52±0.17	0.82±0.19	0.11±0.08	0.09±0.08	0.14±0.10
	- Crowded s.m.	0.54±0.13	0.39±0.14	0.97±0.07	0.10±0.08	0.07±0.06	0.18±0.14
Faster R-CNN [44]	- Single bugs	0.56±0.20	1.00±0.04	0.42±0.21	0.11±0.10	0.21±0.19	0.08±0.08
	- Synth. mixes	0.68±0.16	0.63±0.21	0.83±0.22	0.16±0.14	0.15±0.14	0.19±0.18
	- Crowded s.m.	0.27±0.17	0.81±0.24	0.18±0.16	0.06±0.09	0.19±0.30	0.04±0.06
nn-Detection [4]	- Single bugs	0.03±0.02	0.02±0.01	0.91±0.15	0.00±0.01	0.00±0.00	0.09±0.16
	- Synth. mixes	0.50±0.14	0.41±0.17	0.80±0.22	0.10±0.06	0.08±0.06	0.16±0.11
	- Crowded s.m.	0.60±0.19	0.77±0.23	0.58±0.27	0.11±0.12	0.15±0.16	0.11±0.12

**Table 2:** Detection results on the training tasks obtained during training on the single bugs and the two synthetic mixtures. For U-Net and Faster R-CNN performance is evaluated on a validation split. Since nnDetection is an ensemble created from multiple train-validation splits, we evaluate on all images in the training task. We show mean average precision (mAP) between predicted and ground truth bounding boxes. Higher is better.


Model	- Trained on	mAP	mAP@50
U-Net [8]	- Single bugs	0.72	0.72
	- Synthetic mixes	0.64	0.77
	- Crowded synth. mixes	0.75	0.86
Faster R-CNN [44]	- Single bugs	0.70	0.70
	- Synthetic mixes	0.02	0.02
	- Crowded synth. mixes	0.05	0.05
nnDetection [4]	- Single bugs	0.10	0.10
	- Synthetic mixes	0.59	0.70
	- Crowded synth. mixes	0.49	0.66

models, these results show that even when the objects do not change appearance, the shift in the domain by changing the context surrounding the objects is sufficient for deep learning-based detection models to lose performance.

In preparing the data, we have strived to get the scans of individual bugs and bugs in mixtures to have the same appearance. Since bugs are flexible, there will be differences in the pose of the bugs in mixtures and individual bugs. Larvae will sometimes bend differently and bugs with antennae or legs will bend differently. Also, the cotton is visible and can slightly influence the appearance. Despite these differences, the voxel intensity distributions in individual scanned bugs and mixtures are the same, and therefore the BugNIST dataset is a good basis for developing detection methods that can handle context shift. If the mixtures were (partly) labeled for segmentation, BugNIST could also be used to investigate segmentation algorithms under context shift.



**Fig. 5:** Detection results on real mixed scans. The images are maximum-intensity projections of the volumes. Center points are marked with circles, detections are marked with bounding boxes, and color indicates the bug class. A line from a point to the upper left corner of a bounding box shows a match. If the center point and the bounding box have the same color, then the detection and classification are correct. The shown images are the volumes with the highest scores for the different models. For U-Net the results when trained on synthetic mixes, for Faster R-CNN when trained on synthetic mixes and nnDetection when trained on crowded synthetic mixes.



**Fig. 6:** Distributions of performance scores for the three detection methods.

We chose to annotate only the center positions of the bugs. As the bugs are often elongated and at odd angles, axis-aligned bounding boxes would contain significant empty space and overlap between neighboring boxes—thus offering limited benefit over points. This could be solved with *object* aligned bounding boxes or full voxel annotations. However, this would drastically increase the annotation time. We chose to prioritize having a large number of annotations such that the test set can have a large variety of configurations. Furthermore, given that current methods show ample room for improvement (*cf.* Tab. 1), we

find that point annotations combined with the chosen metric are sufficient to measure method performance and improvements of future approaches.

The methods and experiments are chosen as a basic baseline to illustrate the properties of the BugNIST dataset. From the experiments, we learned that data augmentation by composing synthetic mixtures is a feasible approach for improving detection performance. More investigations into packing algorithms and ways to combine the synthetic mixed bugs with other backgrounds can potentially improve performance further. But it is also clear that there is great potential for developing new methods for object detection in volumes. We chose Detection U-Net, nnDetection, and our adaption of Faster R-CNN to 3D volumes because these methods were available, widely used, and are based on complementary ideas. We did not find more recent alternatives for volumetric analysis.

Our results show that Faster R-CNN performs less than the two other methods. While it scores highest in Tab. 1, it is clear from Fig. 5 that the bounding boxes are of much lower quality. Furthermore, it does not score well on the training tasks, except when trained on individual bugs. To ensure the poor performance of Faster R-CNN was not just the result of too few training volumes, we redid the experiments with 1000 synthetic mixes instead of the original 300. Performance remained the same. It is also noteworthy that nnDetection scores overall lower than U-Net on both the training and evaluation tasks, even though it also has access to the segmentation masks for training.

## 7 Conclusion

With our BugNIST benchmark dataset for domain shift in volumetric data, we demonstrate how detecting objects, that have the same appearance in the source and target domain, fails when the objects in the two domains are placed in different contexts. By contexts, we refer to the appearance of the image surrounding the objects. To our knowledge, the BugNIST dataset is the only dataset designed for developing methods for context shift, *i.e.* domain shift where the objects stay the same but the surroundings change. Solutions will enable an automated annotation of data for volumetric object detection, which is an area that has not received much attention. There are, however, a range of applications where object detection is needed, including medical imaging and emerging areas such as micro-CT scanning of materials. The BugNIST dataset can also be used for investigating other problems such as segmentation under domain shift provided reference annotations are created.

## Acknowledgments

We acknowledge the 3D Imaging Center at DTU, the Infrastructure for Quantitative AI-based Tomography (QUAITOM), supported by the Novo Nordisk Foundation (grant number NNF21OC0069766) and STUDIOS: Segmenting Tomograms Using Different Interpretation of Simplicity funded by the Villum Foundation (grant number VIL50425).

## References

1. Armato, S.G., Huisman, H., Drukker, K., Hadjiiski, L., Kirby, J.S., Petrick, N., Redmond, G., Giger, M.L., Cha, K., Mamonov, A., et al.: PROSTATEx challenges for computerized classification of prostate lesions from multiparametric magnetic resonance images. *Journal of Medical Imaging* **5**(4), 044501–044501 (2018)
2. Armato, S.G., McLennan, G., Bidaut, L., McNitt-Gray, M.F., Meyer, C.R., Reeves, A.P., Zhao, B., Aberle, D.R., Henschke, C.I., Hoffman, E.A., et al.: The lung image database consortium (LIDC) and image database resource initiative (IDRI): a completed reference database of lung nodules on CT scans. *Medical Physics* **38**(2), 915–931 (2011)
3. Baid, U., Ghodasara, S., Mohan, S., Bilello, M., Calabrese, E., Colak, E., Farahani, K., Kalpathy-Cramer, J., Kitamura, F.C., Pati, S., et al.: The RSNA-ASNR-MICCAI BraTS 2021 benchmark on brain tumor segmentation and radiogenomic classification. *arXiv preprint arXiv:2107.02314* (2021)
4. Baumgartner, M., Jäger, P.F., Isensee, F., Maier-Hein, K.H.: nnDetection: a self-configuring method for medical object detection. In: *Proceedings of the Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)*. pp. 530–539. Springer (2021)
5. Cardoso, M.J., Li, W., Brown, R., Ma, N., Kerfoot, E., Wang, Y., Murrey, B., Myronenko, A., Zhao, C., Yang, D., et al.: MONAI: An open-source framework for deep learning in healthcare. *arXiv preprint arXiv:2211.02701* (2022)
6. Charytanowicz, M., Kulczycki, P., Kowalski, P.A., Łukasik, S., Czabak-Garbacz, R.: An evaluation of utilizing geometric features for wheat grain classification using X-ray images. *Computers and Electronics in Agriculture* **144**, 260–268 (2018)
7. Chaudhary, S., Sadbhawna, S., Jakhetiya, V., Subudhi, B.N., Baid, U., Guntuku, S.C.: Detecting COVID-19 and community acquired pneumonia using chest CT scan images with deep learning. In: *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. pp. 8583–8587. IEEE (2021)
8. Çiçek, Ö., Abdulkadir, A., Lienkamp, S.S., Brox, T., Ronneberger, O.: 3D U-Net: learning dense volumetric segmentation from sparse annotation. In: *Proceedings of the Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)*. pp. 424–432. Springer (2016)
9. Cordts, M., Omran, M., Ramos, S., Rehfeld, T., Enzweiler, M., Benenson, R., Franke, U., Roth, S., Schiele, B.: The cityscapes dataset for semantic urban scene understanding. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 3213–3223 (2016)
10. De Carlo, F., Gürsoy, D., Ching, D.J., Batenburg, K.J., Ludwig, W., Mancini, L., Marone, F., Mokso, R., Pelt, D.M., Sijbers, J., et al.: Tomobank: a tomographic data repository for computational x-ray science. *Measurement Science and Technology* **29**(3), 034004 (2018)
11. Dorent, R., Kujawa, A., Ivory, M., Bakas, S., Rieke, N., Joutard, S., Glocker, B., Cardoso, J., Modat, M., Batmanghelich, K., et al.: Crossmoda 2021 challenge: Benchmark of cross-modality domain adaptation techniques for vestibular schwannoma and cochlea segmentation. *Medical Image Analysis* **83**, 102628 (2023)
12. Du, Z., Hu, Y., Ali Buttar, N., Mahmood, A.: X-ray computed tomography for quality inspection of agricultural products: A review. *Food Science & Nutrition* **7**(10), 3146–3160 (2019)

13. Einarsdóttir, H., Emerson, M.J., Clemmensen, L.H., Scherer, K., Willer, K., Bech, M., Larsen, R., Ersbøll, B.K., Pfeiffer, F.: Novelty detection of foreign objects in food using multi-modal X-ray imaging. *Food Control* **67**, 39–47 (2016)
14. European Organization For Nuclear Research, OpenAIRE: Zenodo (2013). <https://doi.org/10.25495/7GXX-RD71>, <https://www.zenodo.org/>
15. Flitton, G., Mouton, A., Breckon, T.P.: Object classification in 3D baggage security computed tomography imagery using visual codebooks. *Pattern Recognition* **48**(8), 2489–2499 (2015)
16. Girshick, R.: Fast R-CNN. In: *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*. pp. 1440–1448 (2015)
17. Girshick, R., Donahue, J., Darrell, T., Malik, J.: Rich feature hierarchies for accurate object detection and semantic segmentation. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 580–587 (2014)
18. Guan, H., Liu, M.: Domain adaptation for medical image analysis: a survey. *IEEE Transactions on Biomedical Engineering* **69**(3), 1173–1185 (2021)
19. Harmon, S.A., Sanford, T.H., Xu, S., Turkbey, E.B., Roth, H., Xu, Z., Yang, D., Myronenko, A., Anderson, V., Amalou, A., et al.: Artificial intelligence for the detection of COVID-19 pneumonia on chest CT using multinational datasets. *Nature Communications* **11**(1), 4080 (2020)
20. Hipsley, C.A., Aguilar, R., Black, J.R., Hocknull, S.A.: High-throughput microCT scanning of small specimens: preparation, packing, parameters and post-processing. *Scientific Reports* **10**(1), 13863 (2020)
21. Isensee, F., Jaeger, P.F., Kohl, S.A., Petersen, J., Maier-Hein, K.H.: nnU-Net: a self-configuring method for deep learning-based biomedical image segmentation. *Nature Methods* **18**(2), 203–211 (2021)
22. Jaeger, P.F., Kohl, S.A., Bickelhaupt, S., Isensee, F., Kuder, T.A., Schlemmer, H.P., Maier-Hein, K.H.: Retina U-Net: Embarrassingly simple exploitation of segmentation supervision for medical object detection. In: *Machine Learning for Health Workshop*. pp. 171–183. PMLR (2020)
23. Jain, Y., Borner, K., Walsh, C., Ruschman, N., Lee, P.D., Weber, G.M., Holbrook, R., Howard, A.: SenNet + HOA - hacking the human vasculature in 3D (2023), <https://kaggle.com/competitions/blood-vessel-segmentation>
24. Jarolmasjed, S., Espinoza, C.Z., Sankaran, S., Khot, L.R.: Postharvest bitter pit detection and progression evaluation in ‘Honeycrisp’ apples using computed tomography images. *Postharvest Biology and Technology* **118**, 35–42 (2016)
25. Jeppesen, N., Christensen, A.N., Dahl, V.A., Dahl, A.B.: Sparse layered graphs for multi-object segmentation. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 12777–12785 (2020)
26. Jin, L., Yang, J., Kuang, K., Ni, B., Gao, Y., Sun, Y., Gao, P., Ma, W., Tan, M., Kang, H., et al.: Deep-learning-assisted detection and segmentation of rib fractures from CT scans: Development and validation of FracNet. *EBioMedicine* **62**, 103106 (2020)
27. Ker, J., Singh, S.P., Bai, Y., Rao, J., Lim, T., Wang, L.: Image thresholding improves 3-dimensional convolutional neural network diagnosis of different acute brain hemorrhages on computed tomography scans. *Sensors* **19**(9), 2167 (2019)
28. Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., Xiao, T., Whitehead, S., Berg, A.C., Lo, W.Y., et al.: Segment anything. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 4015–4026 (2023)



29. LeCun, Y., Bottou, L., Bengio, Y., Haffner, P.: Gradient-based learning applied to document recognition. *Proceedings of the IEEE* **86**(11), 2278–2324 (1998)
30. Li, D., Yang, Y., Song, Y.Z., Hospedales, T.M.: Deeper, broader and artier domain generalization. In: *Proceedings of the IEEE International Conference on Computer Vision*. pp. 5542–5550 (2017)
31. Li, Y., Xie, S., Chen, X., Dollar, P., He, K., Girshick, R.: Benchmarking detection transfer learning with vision transformers. *arXiv preprint arXiv:2111.11429* (2021)
32. Li, Y., Fan, Y.: Medical image segmentation with domain adaptation: A survey. *arXiv preprint arXiv:2311.01702* (2023)
33. Lin, T.Y., Goyal, P., Girshick, R., He, K., Dollár, P.: Focal loss for dense object detection. In: *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*. pp. 2980–2988 (2017)
34. Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft coco: Common objects in context. In: *Proceedings of the European Conference on Computer Vision (ECCV)*. pp. 740–755. Springer (2014)
35. Ma, J., He, Y., Li, F., Han, L., You, C., Wang, B.: Segment anything in medical images. *Nature Communications* **15**(1), 654 (2024)
36. Maier-Hein, L., Eisenmann, M., Reinke, A., Onogur, S., Stankovic, M., Scholz, P., Arbel, T., Bogunovic, H., Bradley, A.P., Carass, A., et al.: Why rankings of biomedical image analysis competitions should be interpreted with care. *Nature Communications* **9**(1), 5217 (2018)
37. Muandet, K., Balduzzi, D., Schölkopf, B.: Domain generalization via invariant feature representation. In: *International Conference on Machine Learning*. pp. 10–18. PMLR (2013)
38. Munkres, J.: Algorithms for the assignment and transportation problems. *Journal of the Society for Industrial and Applied Mathematics* **5**(1), 32–38 (1957)
39. OECD: Health at a Glance 2023 (2023). <https://doi.org/10.1787/7a7afb35-en>
40. Pan, S.J., Tsang, I.W., Kwok, J.T., Yang, Q.: Domain adaptation via transfer component analysis. *IEEE Transactions on Neural Networks* **22**(2), 199–210 (2010)
41. Peng, X., Bai, Q., Xia, X., Huang, Z., Saenko, K., Wang, B.: Moment matching for multi-source domain adaptation. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 1406–1415 (2019)
42. Rawson, S.D., Maksimcuka, J., Withers, P.J., Cartmell, S.H.: X-ray computed tomography in life sciences. *BMC biology* **18**(1), 1–15 (2020)
43. Redmon, J., Divvala, S., Girshick, R., Farhadi, A.: You only look once: Unified, real-time object detection. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 779–788 (2016)
44. Ren, S., He, K., Girshick, R., Sun, J.: Faster R-CNN: Towards real-time object detection with region proposal networks. *Advances in neural information processing systems (NeurIPS)* **28** (2015)
45. Richter, S.R., Vineet, V., Roth, S., Koltun, V.: Playing for data: Ground truth from computer games. In: *Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part II* 14. pp. 102–118. Springer (2016)
46. Ronneberger, O., Fischer, P., Brox, T.: U-Net: Convolutional networks for biomedical image segmentation. In: *Proceedings of the Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)*. pp. 234–241. Springer (2015)
47. Ros, G., Sellart, L., Materzynska, J., Vazquez, D., Lopez, A.M.: The synthia dataset: A large collection of synthetic images for semantic segmentation of urban

- scenes. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 3234–3243 (2016)
48. Sankaranarayanan, S., Balaji, Y., Jain, A., Lim, S.N., Chellappa, R.: Learning from synthetic data: Addressing domain shift for semantic segmentation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 3752–3761 (2018)
  49. Serte, S., Demirel, H.: Deep learning for diagnosis of COVID-19 using 3D CT scans. *Computers in biology and medicine* **132**, 104306 (2021)
  50. Setio, A.A.A., Traverso, A., De Bel, T., Berens, M.S., Van Den Bogaard, C., Cerello, P., Chen, H., Dou, Q., Fantacci, M.E., Geurts, B., et al.: Validation, comparison, and combination of algorithms for automatic detection of pulmonary nodules in computed tomography images: the luna16 challenge. *Medical Image Analysis* **42**, 1–13 (2017)
  51. Simpson, A.L., Antonelli, M., Bakas, S., Bilello, M., Farahani, K., van Ginneken, B., Kopp-Schneider, A., Landman, B.A., Litjens, G., Menze, B.H., Ronneberger, O., Summers, R.M., Bilic, P., Christ, P.F., Do, R.K.G., Gollub, M., Golia-Pernicka, J., Heckers, S., Jarnagin, W.R., McHugo, M., Napel, S., Vorontsov, E., Maier-Hein, L., Cardoso, M.J.: A large annotated medical image dataset for the development and evaluation of segmentation algorithms. *CoRR* **abs/1902.09063** (2019), <http://arxiv.org/abs/1902.09063>
  52. Sun, T., Segu, M., Postels, J., Wang, Y., Van Gool, L., Schiele, B., Tombari, F., Yu, F.: Shift: A synthetic driving dataset for continuous multi-task domain adaptation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 21371–21382 (2022)
  53. Timmins, K.M., van der Schaaf, I.C., Bennink, E., Ruigrok, Y.M., An, X., Baumgartner, M., Bourdon, P., De Feo, R., Di Noto, T., Dubost, F., et al.: Comparing methods of detecting and segmenting unruptured intracranial aneurysms on TOF-MRAS: The ADAM challenge. *Neuroimage* **238**, 118216 (2021)
  54. TorchVision maintainers and contributors: Torchvision: Pytorch’s computer vision library. <https://github.com/pytorch/vision> (2016)
  55. Trusler, M.M., Sturrock, C.J., Vane, C.H., Cook, S., Lomax, B.H.: X-ray computed tomography: A novel non-invasive approach for the detection of microplastics in sediments? *Marine Pollution Bulletin* **194**, 115350 (2023)
  56. Velayudhan, D., Hassan, T., Damiani, E., Werghi, N.: Recent advances in baggage threat detection: A comprehensive and systematic survey. *ACM Computing Surveys* **55**(8) (dec 2022). <https://doi.org/10.1145/3549932>, <https://doi.org/10.1145/3549932>
  57. Venkateswara, H., Eusebio, J., Chakraborty, S., Panchanathan, S.: Deep hashing network for unsupervised domain adaptation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 5018–5027 (2017)
  58. Walsh, C., Tafforeau, P., Wagner, W., Jafree, D., Bellier, A., Werlein, C., Kühnel, M., Boller, E., Walker-Samuel, S., Robertus, J., et al.: Imaging intact human organs with local resolution of cellular structures using hierarchical phase-contrast tomography. *Nature Methods* **18**(12), 1532–1541 (2021)
  59. Wang, M., Deng, W.: Deep visual domain adaptation: A survey. *Neurocomputing* **312**, 135–153 (2018)
  60. Wang, M., Liu, Y., Yuan, J., Wang, S., Wang, Z., Wang, W.: Inter-class and inter-domain semantic augmentation for domain generalization. *IEEE Transactions on Image Processing* (2024)

61. Wang, Q., Bhowmik, N., Breckon, T.P.: Multi-class 3D object detection within volumetric 3D computed tomography baggage security screening imagery. In: Proceedings of the IEEE International Conference on Machine Learning and Applications (ICMLA). pp. 13–18. IEEE (2020)
62. Wang, Q., Breckon, T.P.: Contraband materials detection within volumetric 3D computed tomography baggage security screening imagery. In: Proceedings of the IEEE International Conference on Machine Learning and Applications (ICMLA). pp. 75–82. IEEE (2021)
63. Withers, P.J., Bouman, C., Carmignato, S., Cnudde, V., Grimaldi, D., Hagen, C.K., Maire, E., Manley, M., Du Plessis, A., Stock, S.R.: X-ray computed tomography. *Nature Reviews Methods Primers* **1**(1), 18 (2021)
64. Wu, D., Wu, D., Feng, H., Duan, L., Dai, G., Liu, X., Wang, K., Yang, P., Chen, G., Gay, A.P., et al.: A deep learning-integrated micro-ct image analysis pipeline for quantifying rice lodging resistance-related traits. *Plant Communications* **2**(2) (2021)
65. Xiao, Z., Liu, B., Geng, L., Zhang, F., Liu, Y.: Segmentation of lung nodules using improved 3d-unet neural network. *Symmetry* **12**(11), 1787 (2020)
66. Xie, Y., Ji, Q.: A new efficient ellipse detection method. In: Proceedings of the International Conference on Pattern Recognition (ICPR). vol. 2, pp. 957–960. IEEE (2002)
67. Xu, F., Tenenbaum, J.B.: Word learning as bayesian inference. *Psychological review* **114**(2), 245 (2007)
68. Yan, Y.T., Chua, S., DeCarlo, T.M., Kempf, P., Morgan, K.M., Switzer, A.D.: Core-ct: A matlab application for the quantitative analysis of sediment and coral cores from X-ray computed tomography (ct). *Computers & Geosciences* **156**, 104871 (2021)
69. Zhao, W.X., Zhou, K., Li, J., Tang, T., Wang, X., Hou, Y., Min, Y., Zhang, B., Zhang, J., Dong, Z., et al.: A survey of large language models. *arXiv preprint arXiv:2303.18223* (2023)
70. Zhou, K., Yang, Y., Hospedales, T., Xiang, T.: Deep domain-adversarial image generation for domain generalisation. *Proceedings of the AAAI Conference on Artificial Intelligence* **34**(07), 13025–13032 (April 2020). <https://doi.org/10.1609/aaai.v34i07.7003>, <https://ojs.aaai.org/index.php/AAAI/article/view/7003>
71. Zhou, S.K., Greenspan, H., Davatzikos, C., Duncan, J.S., Van Ginneken, B., Madabhushi, A., Prince, J.L., Rueckert, D., Summers, R.M.: A review of deep learning in medical imaging: Imaging traits, technology trends, case studies with progress highlights, and future promises. *Proceedings of the IEEE* **109**(5), 820–838 (2021)