



Universiteit
Leiden

The Netherlands

Spectral imaging and tomographic reconstruction methods for industrial applications

Zeegers, M.T.

Citation

Zeegers, M. T. (2023, May 31). *Spectral imaging and tomographic reconstruction methods for industrial applications*. Retrieved from <https://hdl.handle.net/1887/3619550>

Version: Publisher's Version

License: [Licence agreement concerning inclusion of doctoral thesis in the Institutional Repository of the University of Leiden](#)

Downloaded from: <https://hdl.handle.net/1887/3619550>

Note: To cite this publication please use the final published version (if applicable).

2

Tomographic workflows for deep learning training data generation

In food inspection, identifying foreign objects is essential. X-ray imaging and machine learning allow for doing this in automated manner. Nevertheless, for reliable performance a substantial body of training data is needed, which is difficult to come by.

In this chapter, we present a Computed Tomography (CT) based method for producing training data for supervised learning of foreign object detection, with minimal labor requirements. In this approach, a few representative objects are CT scanned and reconstructed in 3D. The radiographs that are acquired as part of the CT-scan data serve as input for the machine learning method. High-quality ground truth locations of the foreign objects are obtained through accurate 3D reconstructions and segmentations. Using these segmented volumes, corresponding 2D segmentations are obtained by creating virtual projections. We outline the benefits of objectively and reproducibly generating training data in this way. In addition, we show how the detection accuracy depends on the number of objects used for the CT reconstructions. We show that in this workflow generally only a relatively small number of representative objects (i.e. fewer than 10) are needed to achieve adequate detection performance in an industrial setting.

This chapter is based on:

M. T. Zeegers, T. van Leeuwen, D. M. Pelt, S. B. Coban, R. van Liere, and K. J. Batenburg. “A tomographic workflow to enable deep learning for X-ray based foreign object detection”. *Expert Systems with Applications* 206 (2022), p. 117768.

2.1 Introduction

Foreign object detection in an industrial high-throughput setting is essential for guaranteeing quality and safety of objects processed in factory lines. Foreign objects may, for example, appear in products such as meat, fish or vegetables as small pieces of glass, bones, plastic, wood or stone that could harm consumers [16, 294, 327]. Conventional nondestructive methods for detecting foreign objects include ultrasound imaging, X-ray imaging, magnetic resonance imaging, fluorescence imaging, (hyperspectral) spectroscopic imaging and thermal imaging [116, 162, 200, 207, 210, 303]. X-ray imaging provides the unique opportunity to visualize the interior structure of an object in a fast, low-cost, and non-invasive manner. This enables *X-ray based foreign object detection*, in which the goal is to detect unwanted smaller objects inside base objects based on their distinct attenuation or attenuation patterns, as observed in generated *radiographs* (i.e. standard 2D X-ray images). The possibility to reveal hidden foreign objects in radiographs has led to its extensive use in various industrial applications [80, 111, 153, 186, 191, 207, 323], for which low-cost, adaptive and efficient image processing methods are essential [186, 303]. One way to achieve better discrimination of foreign objects in radiographs is to use multispectral X-ray imaging detectors, simultaneously capturing radiographs at two or more energy levels [199, 274]. As the attenuation properties of each material have their own characteristic dependence on the X-ray energy, these multispectral images can be analyzed to extract material composition information.

However, superposition of materials gives rise to similar levels of intensities for different objects in 2D radiographs. This problem limits the application of commonly used segmentation methods, such as threshold-based, clustering-based, and boundary-based or edge-based segmentation [256, 260], to extract different components of the object. Additionally, high-throughput acquisition may lead to high noise levels in radiographs, and this increases the difficulty of successful foreign object detection even further [186, 303]. Commonly used segmentation methods can be unsuitable in case of poor image qualities caused by conditions such as noise, low contrast and homogeneity in regions close to foreign objects [260]. Most conventional unsupervised methods can therefore not achieve high accuracies [260] without extensive manual parameter tuning to use a method for a specific problem [239, 245].

Machine learning is a powerful tool for recognizing patterns in images [322] and can potentially detect foreign objects in radiographs [327]. Recent machine learning methods address a wide variety of segmentation problems [98, 260], and provide a remarkable improvement over more classical segmentation methods in many practical applications [108]. A key obstacle in the application of machine learning is the need for large datasets [57, 69, 297], which is particularly prominent in machine learning for foreign object detection as each new combination of sample, foreign object, and imaging settings requires additional data. On top of that, supervised

learning uses labelled datasets for training. However, manual annotation (as in e.g. [245, 260]) requires tremendous efforts [8], is time consuming and tedious [276], is subjective and can be prone to errors.

The key contribution of this chapter is the proposal of a workflow based on 3D Computed Tomography (CT) for efficiently creating large training datasets, overcoming the aforementioned obstacle. CT scans of a relatively small number of objects are carried out with low exposure time – as in a high-throughput setting – yielding a large number of radiographs that are used as input for the supervised machine learning method. The same set of radiographs is also used offline for generating multiple high-quality tomographic 3D reconstructions, from which foreign objects can easily be segmented in 3D and projected back onto a virtual 2D detector to give the corresponding ground truth locations of the foreign objects in the radiographs. Without the effort of extensive manual labelling, this results in a large dataset with which deep learning can be carried out to detect foreign objects from fast-acquisition radiographs at a high rate. The example in Figure 2.1 illustrates the difference in ease of segmentation for a CT reconstructed 3D volume versus a 2D radiograph. Whereas segmenting the foreign object in a radiograph is a challenging task, simple global thresholding can be applied to the CT volume to separate the foreign object from the base object. Additionally, more sophisticated and accurate segmentation and denoising rules can be imposed on 3D volumes [98, 173, 217] than on 2D radiographs.

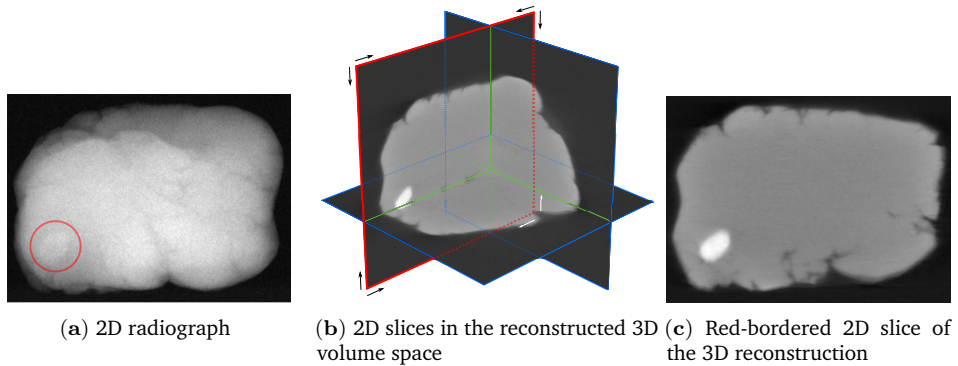


Figure 2.1: Different views of an imaged product (Play-Doh) with a foreign object inserted (a piece of gravel). A 2D radiograph with the location of the foreign object (red circle) is shown (a), as well as multiple slices through the 3D volume of the reconstructed object (b), of which the slice with the red border is highlighted (c). The images show the difference in contrast: the foreign object is much easier to distinguish based on intensity values in the reconstructed 3D volume (b and c) than in the 2D radiograph (a).

The structure of the chapter is as follows. Section 2.2 provides the background of applying machine learning for foreign object detection, and Section 2.3 explains the proposed method of data generation to apply machine learning. In Section 2.4, the workflow is demonstrated in a laboratory experiment, and shows how the number of imaged objects affects the detection accuracies. Additionally, the robustness of the workflow is analyzed. Section 2.5 discusses various aspects of the results and the flexibility and modularity of the workflow. Section 2.6 presents the conclusions from this work.

2.2 Preliminaries

In this section, we introduce the X-ray foreign object detection problem and the machine learning concepts used in this chapter.

2.2.1 Foreign object detection with X-ray imaging

We consider the problem of foreign object detection in an industrial high-throughput conveyor belt setting. The problem and the usage of X-ray imaging to solve this are schematically shown in Figure 2.2. In foreign object detection, the aim is to correctly determine for each object whether a foreign object is contained in it or not, for instance a piece of bone within a meat sample.

For this problem, we focus on finding an accurate *segmentation* for each radiograph. A segmentation partitions an image into sets of pixels with the same label. In our case, the formed segmented image is binary and indicates on which detector pixels a foreign object is projected. The segmentation depends on the type of objects that are considered to be foreign (by for instance a manufacturer). Any further classification (based on the minimum size of a foreign object for example) can be carried out after the segmented image is produced.




Figure 2.2: A schematic overview of the foreign object detection problem and the segmentation-based approach to solving it. Each object (a) is assumed to have a correct segmentation (e). By using X-ray imaging (b), a radiograph of the object (c) can be acquired. Using a segmentation method (d), a segmented image (e) can be produced. The main challenge is to find a suitable segmentation method such that this approach to foreign object detection produces the correct results.

Throughout this chapter, we use the term *radiograph* for radiographs corrected using flatfield radiographs (without an object) and darkfield images (without the X-ray beam) that serve as input to the segmentation method. The quality of a radiograph depends on a number of properties of the scan, including exposure time, tube intensity, photon energy windows and the geometric setup [242]. In a high-throughput setting, the steps in Figure 2.2 should be fast to carry out, typically resulting in high noise levels and a challenging segmentation task.

2.2.2 Supervised learning

Machine learning is a widely used approach for difficult imaging tasks, as it can extract complicated patterns from complex images. In the foreign object detection problem, supervised machine learning can be used to learn the segmentation task such that it generalizes well for all possible fast-acquisition radiographs of similar objects with similar acquisition settings. To do so, a set of examples $\{(x_i, y_i)\}_{i=1}^N$ is used, where $\{x_i\}_{i=1}^N$ are acquired radiographs and $\{y_i\}_{i=1}^N$ are their corresponding foreign objects segmentations. The aim is to find the unknown segmentation function F that maps each radiograph x_i to its segmentation y_i . To find an approximate solution that generalizes well, the set of images is partitioned into a training set, a validation set and a test set. The training set is used to learn the function F_{train} that minimizes the loss L on the training set, which is the sum of errors between the segmented images $F_{\text{train}}(x_i)$ produced by the segmentation function and the true segmented images y_i . To find a suitable segmentation function, a (convolutional) neural network is often used as a model and parametrized using weights and biases that are optimized during the training process. While carrying out the training with a chosen loss function and optimization algorithm, the performance of the model is evaluated on the validation set. Several stopping criteria can be used for this, for example stopping the training when the error on the validation set increases, or training for a fixed time (and recording the network that gives the best results on the validation set). To avoid any bias towards the training and validation data, the accuracy of the trained model is finally assessed using the test set.

Since the introduction of Fully Convolutional Networks [171], in which successive contracting convolutional layers are utilized for pixel-wise semantic segmentation, many convolutional neural network (CNN) architectures have been proposed that can be used for the object segmentation task. U-Net changes the FCN architecture by - along with downsampling operators and skip connections - introducing up-sampling operators instead of pooling operators, giving it an U-shaped appearance [240]. Similarly, Deconvnet [212] also introduces an auto-encoder structure with deconvolution and unpooling operations (without skip connections). The success of these methods on medical image segmentation and object detection spawned other commonly used CNN architectures for segmentation such as SegNet [22], RefineNet [167], PSPNet [321], and Mask R-CNN [115] for instance segmentation. Although

some of the listed architectures need relatively few training examples for successful segmentation, the annotation of these examples still requires considerable efforts.

2.3 Proposed method for training data acquisition

When attempting to perform machine learning for X-ray based foreign object detection, the major obstacle is to acquire (manually) annotated training data. In this section, we explain the methodology of our proposed CT-based workflow for efficiently creating this annotated training data.

Our proposed workflow for using CT to obtain annotated training images is schematically displayed in Figure 2.3. First, we select a set of representative objects as training objects (Fig. 2.3a). For each object, a set of fast-acquisition radiographs is collected from a set of predefined angles (Fig. 2.3b). These fast-acquisition radiographs will form the input set of the intended training dataset (Fig. 2.3c). The total number of examples in the resulting dataset is the number of training objects multiplied by the number of selected angles.

The same set of radiographs is used to carry out a tomographic reconstruction of the object and acquire high-quality CT volumetric data (Fig. 2.3d and e). The next step is to segment the reconstructed volume such that a possible foreign object is separated from the base object (Fig. 2.3f). This segmentation step can be automated and many methods are available to implement this [158]. Here, we consider volumetric segmentation methods that consist of a global thresholding step.




Figure 2.3: The complete workflow of data acquisition (a,b) and the generation of training data (c,h) for deep learning driven foreign object detection, through 3D reconstruction from the CT scan (d, e), segmentation (f), and virtual projections (g). The reconstruction reveals the hidden foreign objects inside the main object. Note that the projection data (d) are usually just the set of fast-acquisition radiographs (c).

Binary segmentation by global thresholding is defined by the following function $S : \mathbb{R} \rightarrow \{0, 1\}$ that acts on every voxel z_{ijk} in reconstruction volume z :

$$S(z_{ijk}) = \begin{cases} 1 & z_{ijk} \geq \theta, \\ 0 & z_{ijk} < \theta. \end{cases}$$

Here, θ is the segmentation threshold. The more angles and other high-quality settings are used to obtain projection data, the easier it is to accurately segment the foreign object. Easier segmentation can also be accomplished by carrying out a separate high-quality scan of the same object and making a reconstruction with these high-quality radiographs. Additionally, for segmentation, prior information about the objects can be used, such as bounding boxes on the foreign object location [144]. Also, 3D denoising [70, 119] can be used to remove non-foreign object pixels captured by the thresholding operation.

From the constructed foreign object segmentation, virtual ground truth projections are generated by simulating projections of the foreign objects onto a virtual detector (Fig. 2.3g). This results in the set of ground truth images, which will serve as target images in the machine learning procedure (Fig. 2.3h). These virtual projections need to be taken under the same angles as in the fast-acquisition scan (Fig. 2.3b). When this procedure is repeated for all objects, this results in a large dataset with annotated training examples with which supervised machine learning can be carried out (Fig. 2.3c and f). The trained model can then be applied to similar new objects scanned in the same fast-acquisition setting, without the need for acquisition of high-quality radiographs or CT scans.

2.4 Experiments and results

In this section, we demonstrate the proposed workflow using the in-house FleX-ray CT system at CWI [62] (Fig. 2.4), and investigate the relation between machine learning performance and the number of training objects used.

2.4.1 Base objects and foreign objects

As test objects, we use base objects that are created from a fixed amount of modelling clay (Play-Doh, Hasbro, RI, USA). Play-Doh is primarily made of a mixture of water, salt and flour and we therefore consider it to be a representative example of products in the food industry, where foreign objects may be pieces of stone, plastic, or metal. A basic shape is deformed and remolded for every object instance (Fig. 2.5a) in such a way that they are similar from object to object, but still exhibit some natural variation. For the foreign objects we choose to use gravel (Fig. 2.5b), with the stones having an average diameter of ca. 7mm (ranging from 3mm to 11mm). These stones have slight variations in shape and material. We





Figure 2.4: The scanning setup in the FleX-ray laboratory with the X-ray source on the left and the detector on the right.



(a) Play-Doh

(b) Stones

(c) Play-Doh with stone inserted

Figure 2.5: An example of a base object (a) and examples of foreign objects (b) used in the laboratory experiments, as well as an example of a combined object (c).

create 3 objects with three inserted stones, 35 with two stones, 62 with one stone (Fig. 2.5c) and 11 without a stone.

2.4.2 CT scanning and data preparation

A fast CT scan is made for each of the objects, which yields both a series of radiographs (i.e. the X-ray projections) and a reconstructed 3D volume of the object. The objects are scanned in the FleX-ray laboratory [62] (Fig. 2.4). The FleX-ray CT scanner has a cone-beam microfocus X-ray point source with a focal spot size of $17\ \mu\text{m}$, and a DEXELA1512NDT detector. The source, object and detector positions can be configured flexibly, and are arranged such that the distance between the source and detector is 69.80 cm, and the distance between the source and the object 44.14 cm. For the radiographs a voltage of 90 kV with a power of 20 W is used, while the exposure time is kept low at 20 ms, with the intention to emulate the imaging conditions of in-line industrial systems and produce sufficiently noisy radiographs. To achieve high-quality reconstructions, 1800 projections of each object are obtained at equidistant angles over a full 360° rotation. All projection angles are precisely recorded during the scan for the later stages of the workflow. Before and after each scan, 10 darkfield images and 10 flatfield projections are obtained. Each object is positioned in a random manner, and the cylinders may therefore be standing upright or be laying down on the long edge. Example radiographs are shown in Figure 2.6. Separating the projected foreign objects from the base object in these radiographs is not a trivial task, illustrating the problem of obtaining annotated training data for automated segmentation using machine learning directly from these images.

The Simultaneous Iterative Reconstruction Technique (SIRT) [140, 265] algorithm (100 iterations) as implemented in the ASTRA toolbox [1, 2] is used to compute the reconstructed 3D CT volume of the object. A visualization of the reconstruction from the third object in Figure 2.6 and its foreign object is shown in Figure 2.7. The CT reconstruction allows to slice the object along different axes. As the CT voxel intensity is directly related to the attenuation coefficient of the material in a voxel, the segmentation task for the 3D CT volume is, in this




Figure 2.6: Example of radiographs (size 965×760 pixels) of three objects that are scanned. In the first and second radiographs the foreign object is clearly visible, but in the third it is more difficult to distinguish it from the base object, even though it is visible on the bottom left.




Figure 2.7: Visualization of reconstructions by 3D rendering (a) and slicing (b) of the third object in Figure 2.6.

case, much more straightforward and can be carried out by global thresholding (see Appendix A.1 for additional details on intensity value distributions). Therefore, a simple global threshold based on Otsu’s method [214] is sufficient to segment the foreign objects.

From the 3D segmented objects obtained from the CT scans, 2D segmentations for the individual radiographs are computed. This is done by computing the projections of the segmented parts with the ASTRA toolbox using the same geometric properties and recorded angles as in the radiograph acquisition of the actual CT scan, to ensure geometric consistency in the training examples. Every nonzero pixel on the detector is marked as a projected foreign object location. The result is a dataset containing 1800 radiographs and corresponding segmented images for each object.

2.4.3 Machine learning

We use the Mixed-Scale Dense (MSD) [227] and the common U-Net convolutional neural network architectures [240] to train the task of image segmentation. For our experiments with U-Net, we have slightly changed the architecture, as we observed this improved performance in the experiments compared to the standard version. We downsample twice, with a stride of 2. The initial number of feature maps is set to 128, and the number of feature maps doubles for each downsampling layer. For upsampling, bilinear interpolation is used. A spatial 3×3 convolution operation with zero padding and a ReLU activation function are carried out before and after all downsampling and upsampling operations. The biases and convolution weights are initialized by sampling from $\mathcal{U}(-\sqrt{k}, \sqrt{k})$, with $k = 1/c_{\text{in}} \cdot a^2$ being the range, c_{in} the number of input channels and a the kernel size. ADAM optimization on the average of the binary cross entropy loss and the dice loss [135, 270] between the data and the predictions is used for training. The network is implemented with PyTorch [222, 223]. For comparison between architectures, we also use the

MSD network for training. MSD is a compact network architecture that has been demonstrated to be suitable for real-time segmentation of X-ray and CT images using relatively few training examples compared to larger networks [227], including the U-Net architecture. We use a depth of 100 intermediate layers and width of 1 channel per intermediate layer and increase the dilation parameter repeatedly from 1 to 10 dilations in each layer, which are common settings for the MSD network [154, 226, 227]. Xavier initialization is used for the convolution weights. ADAM optimization [146] is used during training on the cross-entropy loss between the ground truth and the segmented images, and the batch size of training examples is set to 10. We use the GPU implementations in Python that are available [225, 227]. For both architectures, the learning rate is set to 0.001 and all networks are trained on a GeForce GTX TITAN X GPU with CUDA version 10.1.243. All hyperparameters of both network architectures are kept the same during all experiments. Data augmentation is applied by rotation and flipping of the input examples. All networks are trained for 9 hours, and the network with parameters resulting in the lowest error on the validation set is used for testing.

With these networks, we carry out an image-to-image training from radiographs (Fig. 2.3c) to their corresponding foreign objects segmentations (Fig. 2.3f). For training, 60 randomly chosen base objects containing a foreign object are used. The remaining 51 objects are used for testing. All images are resized using cubic interpolation to 128×128 to speed up the training process (global thresholding with parameter $\theta = 0.5$ is applied to the resized ground truth images to make these binary again). We test the performance of the trained networks for different numbers of objects included in the training scheme. To compare the workflow with labor-intensive 2D data annotation, we compare the following training strategies:

- **Workflow approach:** For each network, we fix the total number of training examples to 1800. A random but fixed order of the 60 training objects is created and the first i objects among these are used for the training set. The training examples are selected from the set of radiographs and ground truths created by the workflow from these i training objects in equal amounts. Every 10th example is used for validation during training.
- **Classical approach:** For each network, only one randomly chosen training radiograph with the corresponding ground truth (generated using the workflow) is selected for each of the first i included training objects. The resulting set of training examples is separated such that $9/10$ part is used for training (rounded down to the nearest integer) and $1/10$ part is used for validation (rounded up).

2.4.4 Quality measures

To evaluate the accuracy of the trained networks on the test set, in which the target images are generated using the workflow on the test objects. Three different measures on the segmented images and the corresponding target images are computed. The collection of these measures both assess the image segmentation accuracy and the object detection accuracy. An image segmentation accuracy is based on the classification of each pixel in the segmented image, and there are standardized ways to measure this that do not depend on any parameters [105]. An object detection accuracy compares connected components (groups of pixels connected by their edges) in the segmented image with the ground truth images. Although these accuracy measures require additional parameters to define the notion of detection, they are more relevant to the foreign object detection application.

The first measure is an *image-based average class accuracy* (also called *balanced accuracy* [105]) to assess the accuracy of a produced segmentation. The average class accuracy of a segmented image relative to the target image is given by the sum of the true positives divided by the true positives and false negatives (the recall) of each class, averaged over the number of classes. In the binary case this becomes

$$\frac{1}{2} \left(\frac{\text{TP}_{\text{FO}}}{\text{TP}_{\text{FO}} + \text{FN}_{\text{FO}}} + \frac{\text{TP}_{\text{BG}}}{\text{TP}_{\text{BG}} + \text{FN}_{\text{BG}}} \right). \quad (2.1)$$

Here, TP_{FO} , FN_{FO} , TP_{BG} and FN_{BG} are the true positives and false negatives of the foreign object and the combined base object and background pixel classifications respectively over the entire segmented image relative to the target image. The average class accuracy as given in (2.1) is averaged over all target images.

The second measure is an *object based detection rate*. A *connected component* is a maximal set of nonzero-valued pixels such that each pixel is reachable from another pixel in the set via a sequence of neighboring pixels in the set. Each connected component in the target image with a minimum size of 8 pixels (0.05% of the image size) is considered as an object that should be detected. We define such an object as *detected* if its pixel-wise recall relative to the segmented image is higher than a certain threshold η :

$$\frac{\text{TP}_{\text{obj}}^{\text{tar}}}{\text{TP}_{\text{obj}}^{\text{tar}} + \text{FN}_{\text{obj}}^{\text{tar}}} > \eta. \quad (2.2)$$

Here, $\text{TP}_{\text{obj}}^{\text{tar}}$ and $\text{FN}_{\text{obj}}^{\text{tar}}$ are the true positive and false negative pixels in the target object relative to the segmented image. The threshold indicates the percentage of pixels of a projected foreign object in the target image that should be indicated as foreign object pixels in the segmented image produced by the network to be marked as a detected object. In our experiments, we set $\eta = 0.3$. We define the *detection rate* as the percentage of components in all target images for which condition (2.2) holds.

The third measure is an *object based false positive detection rate*. Each connected component in the segmented image with a minimum size of 8 pixels is considered as a potentially detected object. We define such a potentially detected object as a *false positive* if its pixel-wise recall relative to the target image is lower than a certain threshold δ :

$$\frac{\text{TP}_{\text{obj}}^{\text{seg}}}{\text{TP}_{\text{obj}}^{\text{seg}} + \text{FN}_{\text{obj}}^{\text{seg}}} < \delta. \quad (2.3)$$

Here, $\text{TP}_{\text{obj}}^{\text{seg}}$ and $\text{FN}_{\text{obj}}^{\text{seg}}$ are the true positive and false negative pixels in the segmented object relative to the target image. The threshold indicates the percentage of pixels of a foreign object in the segmented image produced by the network that are correctly labelled as foreign objects compared to the foreign object in the target image. In our experiments, we set $\delta = 0.3$. We define the *false positive detection rate* as the percentage of potential objects in all segmented images for which condition (2.3) holds.

2.4.5 Results

For the test set, we select a random angle and an orthogonal one for each test object, making the total number of testing radiographs 102. We measure the average class accuracy, the object based detection rate and the object based false positive detection rate of segmentations created by the network on the projections from the test set. The results are given in Figure 2.8.

For all measures, the quality of the foreign object segmentations in the radiographs using networks trained with the workflow data is low for a few training objects. This initially improves with the addition of relatively few training objects, but this improvement stagnates beyond 20 objects. However, the detection accuracy still shows slight improvements beyond this point, but almost completely stabilizes from 40 objects onwards. Based on a decided accuracy goal, a certain number of objects need to be scanned and used for training to achieve that accuracy. The false positive rate decreases strongly and maintains a low level value from including 3 objects in the training onwards. Note that the results between the U-Net and MSD architectures agree well with each other.

When we compare the usage of a fixed number of training radiographs among all training objects with the classical approach of using only one radiograph per object, we see that this leads to inferior results in all aspects. The average class accuracies and the object based detection rates are lower for all numbers of included training objects, while the false positive rates are higher. The difference between architectures only shows for the false positive detection rate, which is generally higher with the U-Net architecture.

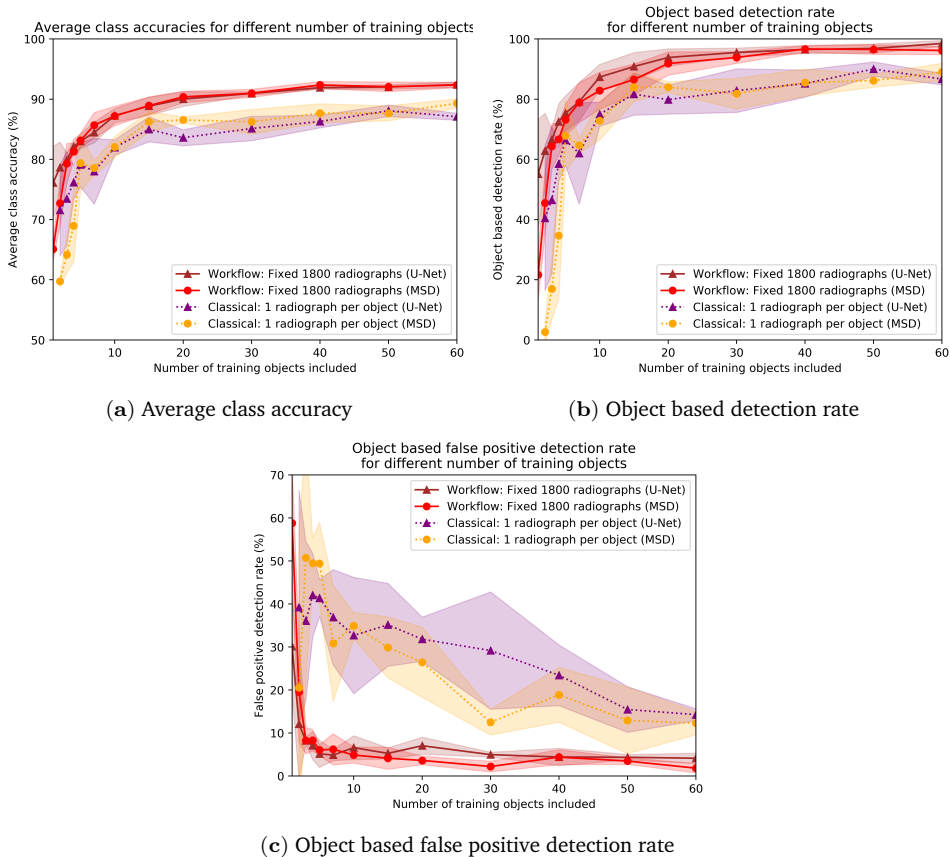



Figure 2.8: The average class accuracy (a), the object based detection rate (b), and the object based false positive rate (c) of segmentations with trained U-Net and MSD networks on laboratory data for different number of training objects. The results are shown for the fixed number of training radiograph approach (workflow) and the one training radiograph per object approach (classical). The results are averaged over 5 trained networks, with a different training object order for each run. The shaded regions indicate the respective standard deviations.

2.4.6 Laboratory experiments with many foreign objects


A natural way to reduce the number of objects used for training that need to be scanned for obtaining accurate segmentations may be to include more foreign object in the imaged objects. To test this, we repeat the experiments of the previous section, but we insert 5 to 8 foreign objects instead of 0 to 2. The foreign objects are placed within the base object such that overlapping of foreign objects in the radiographs is minimized. We have scanned an additional set of 20 objects with these characteristics. An example of a radiograph of an object with many foreign objects is shown in Figure 2.9a. We compare the following training strategies in

which the workflow data come from the following sets of training objects:


- **Few foreign objects:** Base objects with 0 to 2 foreign objects
- **Many foreign objects:** Base objects with 5 to 8 foreign objects
- **Mixed:** 50% – 50% mix of base objects with 0 to 2 foreign objects and base objects with 5 to 8 foreign objects.




(a) Base object with many foreign objects



(b) Average class accuracy



(c) Object based detection rate



(d) Object based false positive detection rate

Figure 2.9: Example of a radiograph of a base object with many foreign objects (a). The eight foreign objects are vertically placed in the object, although there still may be some overlap. The average class accuracy (b), the object based detection rate (c), and the object based false positive detection rate (d) of segmentations with trained U-Net and MSD networks for different number of training objects and training strategies are shown. The results are averaged over 5 trained networks, with a different training object order for each run. The shaded regions indicate the standard deviations.

All networks are evaluated on the testing set from the previous section (with test objects containing few foreign objects). The average class accuracies, detection accuracies and the false positive rates of the trained neural networks with these schemes on the test set are shown in Figure 2.9. From the graphs in Figures 2.9b and 2.9c we see that the average class accuracies and detection accuracies are higher for the many foreign object training scheme, but Figure 2.9d indicates that false positive rate is also roughly 5 times higher. The mixed approach appears to find middle ground between the two other approaches for all measures. We see that from 20 objects onwards the mixed approach is as good as the approach with few foreign objects in terms of the false positive rate, while being superior in terms of average class accuracy and detection accuracy for up to 40 training objects. This shows that including many foreign objects in the training set for detecting few to no foreign objects in the test set has limited additional value, but mixing these with examples with objects containing a few foreign objects may result in higher detection quality while maintaining a similar false positive detection rate.

2.4.7 Robustness of the workflow

In the previous experiments, the trained networks are tested on a set of projections that are generated using the same 3D segmentation threshold parameter in the workflow as in the generation of the data for the training and validation sets. To assess the robustness of the workflow to different segmentation parameters, we generate the training datasets with different values of the segmentation parameter θ (see Figure 2.10a). For each of these values, networks are trained and assessed on the test set from the previous sections. The number of training objects that are included in the workflow is fixed to 10 (which has led to equivalent results in the previous experiments as with 60 objects in the classical approach).

In Figure 2.10, the average class accuracies, detection accuracies and the false positive rates of the trained neural networks are shown for the different thresholds. The results for U-Net and MSD are very similar. As the threshold value increases, the average class accuracy decreases, with significantly lower values for $\theta = 0.014$ and $\theta = 0.015$. The same holds for the detection rate, but it reaches a plateau between $\theta = 0.009$ and $\theta = 0.013$ where this accuracy measure gives similar values. For low values of the threshold parameter, the false positive values are high, and from $\theta = 0.011$ and higher these are low and similar to each other. Taken together, threshold parameters between $\theta = 0.011$ and $\theta = 0.013$ lead to very similar results. We conclude that for the class of objects considered in these experiments, the workflow is robust against moderate variation of the segmentation parameter and that suboptimal segmentation methods can also be used in the workflow.

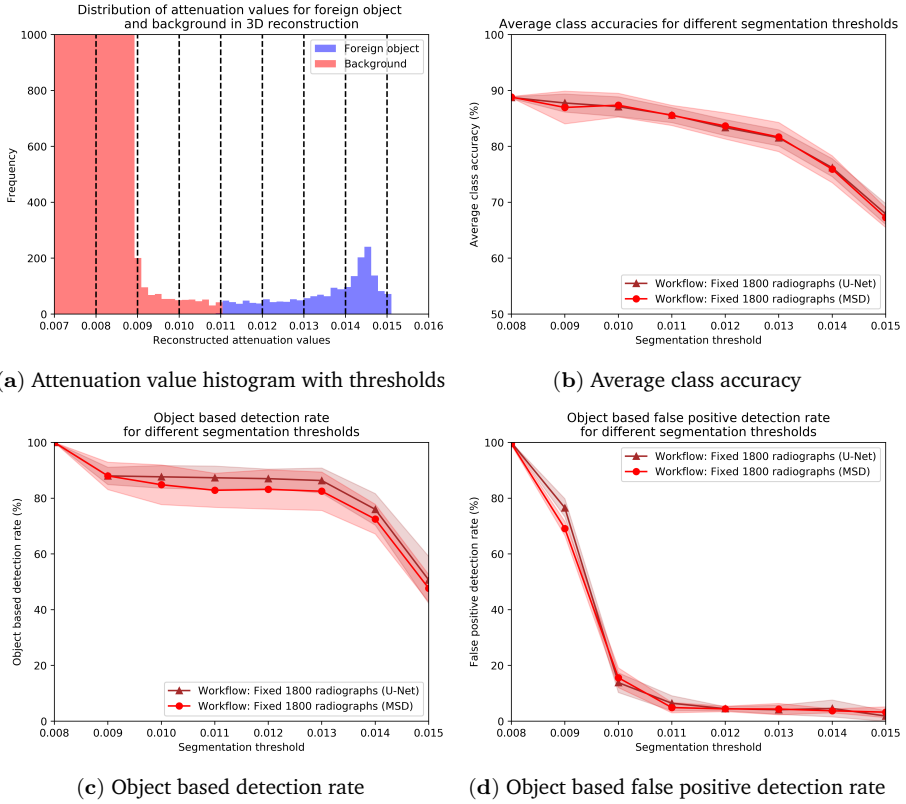


Figure 2.10: The eight considered thresholds for the generation of the training datasets in the workflow, drawn in the histogram of attenuation values of the third object in Figure 2.6, and the average class accuracy (b), the object based detection rate (c), and the object based false positive detection rate (d) of segmentations with trained U-Net and MSD networks on data resulting from these segmentation thresholds. The results are averaged over 5 trained networks, with a different training object order for each run. The shaded regions indicate the standard deviations.

2.4.8 Simulation experiments

In this section, we will demonstrate the workflow in a controlled simulated setting. In this way, we can verify the results with larger training and test sets when more objects are available. Furthermore, the test set previously consisted of data generated with the workflow, but in a simulated setting ‘absolute’ ground truth can be created for the test set by directly projecting the simulated foreign objects (see Figure 2.11). We verify that the proposed workflow (with CT scanning, reconstruction and segmentation) results in segmented foreign objects of which the projections are similar to absolute ground truth projections, which further supports the confidence we can have in the experimental test results.




Figure 2.11: The complete workflow of data acquisition (a,b) and the generation of training data (c,h) for the simulation experiment by 3D reconstruction from the CT scan (d, e), segmentation (f), and virtual projections (g). The 3D reconstruction reveals the hidden foreign objects inside the object. The dotted green arrow (a to f) indicates that because of the simulated nature of the objects, the reconstruction and segmentation steps are skipped for the generation of ground truth for objects in the test set.

We have generated a set of 500 objects, each in an object space of 128^3 voxels. Each object is a cube of size 64^3 voxels, which is placed in the center of the volume. To create sufficient variety among the objects, the cube is cut off by eight planes. For each corner of the cube, a plane is created by selecting points on each of the three outgoing edges of the corner, randomly between the corner point and the midpoint of that edge. The pixels are cut off whose location is on side of the plane opposite to the center of the cube. See Figure 2.12 for a visualization. Additionally, we rotate the resulting object with random angles around all axes. After that we include a foreign object as an ellipsoid with a radius randomly chosen between 3 and 7 voxels at a random location within or on the edge of the base object. These ellipsoids have a random orientation as well. As a result, the foreign objects vary in shape, size, orientation and location. With 50% probability, we include two of these foreign objects instead of one in the base object.




Figure 2.12: The process of cutting off corners with planes from cubes for creating the simulated base objects. The red stripes indicate edge midpoints and the blue dots are the randomly chosen points between those midpoints and the corners.




Figure 2.13: Example radiographs of five simulated objects. Foreign objects are located at various positions in or on the border of the base object in the radiographs, and there can be one or two of these present.

Based on the spectral properties of the assigned materials, we create simulated radiographs (Fig. 2.11b). Details of the computation can be found in Appendix B.2. First, we make projections of each material separately by computing cone beam forward projections using the ASTRA toolbox [1, 2]. From this, the simulated radiographs are computed by taking the spectral properties of each material into account (taken from the National Institute for Standards and Technology (NIST) [130]). We model the foreign objects as bone and the base object as tissue for each object. We take the spectral material characteristics between 15 keV and 90 keV into account, and use an exposure time of 0.002 seconds for each radiograph, for which the Poisson noise that is applied is relatively high. These settings are chosen such that there is sufficient contrast in the radiographs, but not as much that it can be very easily identified with simple segmentation methods. The simulated detector size – and therefore the projection image size – is 128×128 pixels. Examples of radiographs from five objects are given in Figure 2.13.

A total of 100 objects are reserved as training objects, while the other 400 objects are reserved for testing. For each training object, the ground truth corresponding to each radiograph is generated with the workflow, with the same strategy and parameters as in Section 2.4.2. Global thresholding with parameter value $\theta = 0.04$ is used for the reconstructions. For each test object, the ‘absolute’ ground truth corresponding to each radiograph is generated by directly projecting the virtual foreign objects (Fig. 2.11a and f), thereby skipping the reconstruction and segmentation steps. The projections are segmented such that every nonzero pixel on the detector is a projected foreign object location.

To verify that the direct use of the generated 3D volumes results in very similar ground truth projections compared to when the workflow is followed, the resulting ground truth projections are compared for the training set. The Jaccard index between the resulting ground truth pairs, averaged over all projection angles for all 100 training objects, is 0.961 for SIRT with 100 iterations. This result indicates that the resulting ground truth projections resulting from both approaches are very similar, and both are likely to yield the same quality measures.

To further confirm this, the training of networks as described in Section 2.4.3 is repeated with the simulated projections, with the trained networks this time being evaluated on the test set with ‘absolute’ ground truth. The results for the three measures are given in Figure 2.14, and are in accordance with the experiments with the laboratory data. A notable difference is that the average class accuracy and detection accuracy reach their maximum values for a relatively lower number of training objects (and the same goes for the minimum value of the false positive rate). This is most likely because the simulated objects are less complex, resulting in radiographs with less complicated structures. Nevertheless, the results again show inferior results for the approach where one radiograph per training object is used, since 100 objects are needed to reach similar quality measure values as for the workflow with only 7 objects.




Figure 2.14: The average class accuracy, the object based detection rate, and the object based false positive detection rate of segmentations with trained U-Net and MSD networks on simulated data for different number of training objects and different training strategies. The results are averaged over 5 trained networks, with a different training object order for each run. The shaded regions indicate the standard deviations.

2.5 Discussion

Regarding the difficulty of transferring the method to industrial settings, practical issues such as differences in the CT setup can be expected, but we expect that these are straightforward to overcome. In general, we expect that the proposed workflow can be transferred to industrial applications if the following three conditions are met. First of all, the materials should have different absorptions in the energy range of acquired radiographs (determined by the peak voltage of the source and the energetic detection range of the detector) to be distinguishable at all. Secondly, it should be possible to carry out a segmentation in the reconstruction domain, and the foreign objects should therefore not be too small. Lastly, the foreign objects should be detectable in the radiographs.

In our experiments, the Play-Doh is selected to be representative of many example products in the food industry, and the stones for the related foreign objects. This particular case meets the conditions stated above, and therefore stone can be detected in our examples. Moreover, we expect no problems with metal detection because of its higher X-ray attenuation and visibility in the radiographs, and this high attenuation may allow for even smaller metal foreign object sizes as they are more likely to still appear in the reconstruction and radiographs. On the other hand, if the objects contain large metal pieces and other materials need to be detected, it could lead to artifacts in the reconstructions, but there are many artifact reduction methods available that can be used to mitigate this [101]. Regarding other less dense materials such as plastics, successful application of the workflow strongly depends on the visibility of the foreign objects in the radiographs in the first place. If foreign objects are impossible to discern in radiographs, creating training data by manual annotation is also not possible. Of course, many solutions can be proposed for this invisibility problem, but this discussion is independent of the workflow for training data generation. However, even without advanced imaging methods, applying the workflow and generating 2D ground truth could lead to networks retrieving patterns in the radiographs that are difficult to find by human inspection. Additionally, we have shown in the experiments that multiple foreign objects in a sample do not pose a problem. When the material types among these are varied, we do not expect any problem as long as they are distinguishable in the radiographs. Finally, when transferring the workflow to an industrial setup, we may expect practical issues regarding the CT setup, but this is not a fundamental problem of the proposed method.

Overall, the graphs presenting the foreign object detection accuracies of stones in Play-Doh in Section 2.4 indicate an increase of segmentation and detection accuracy with increasing the number of objects from which the training data are created. The accuracies initially increase strongly with the number of training objects but this increase decays when the number of training objects is further increased. The detection rates and false positives rate introduced in Section 2.4.4

depend on the thresholds η and δ , respectively. If a higher threshold η is used in condition (2.2), the detection rate will decrease, because fewer objects in the target images will meet this condition. Conversely, if the threshold decreases, the detection rate will increase. Similarly, if the δ threshold in condition (2.3) is increased, there will be more false positives (and fewer when the threshold is decreased). Changing the thresholds η and δ will therefore change Figures 2.8, 2.9 and 2.14 as well, but we expect the overall shape of the curves and their relative distances to each other will remain similar.

When the thresholds are fixed, the maximum detection accuracy that can be achieved depends on the nature of the foreign detection problem. For instance, if the X-ray flux is low and the noise is high, foreign objects are more difficult to detect from the radiographs. In the case of the laboratory experiments, foreign objects are difficult to detect when the cylindrical shape is located with the long edge on the ground and oriented orthogonal to the detector. The radiographs should contain sufficient discriminatory information such that foreign object detection with deep learning is possible. Additionally, for the dataset to be suitable for supervised machine learning, the ground truth should also be of sufficient quality, although this seemed to be less of an issue in our experiments as we observed no negative effects from occasional noise in the ground truth on the training and detection accuracy.

With the above considerations in mind, the workflow is designed to be modular. Every stage of the proposed workflow can be designed according to the available data-acquisition equipment, the intended detection accuracy, the type of base objects and foreign objects, and the available computer memory, among other things. We highlight some possible considerations for every stage:

- **Objects** (Fig. 2.3a): The set of objects can be enlarged or diversified when the accuracy of the trained neural network is not satisfactory. Also, more objects can be added to obtain a more diverse representation of objects when a more diverse array of objects or orientations are considered to be subjected to X-rays in the industrial application, such as on a conveyor belt. When a completely new type of objects is considered, these objects should be added to the workflow as well.
- **Scanning routine** (Fig. 2.3b): In our experimental setting we have used data resulting from low exposure times as input for both the neural networks and the reconstruction algorithm. If the foreign objects turn out to be too difficult to separate in the reconstructions, more scanning angles may be considered. Additionally, if the factory settings are allowed to be altered, higher fluxes, different tube voltages or longer exposure times can be used to obtain radiographs of higher quality, as long as the processing times remain acceptable. Also, more discrimination can be achieved by applying

spectral imaging (dual-energy [235] or multi-energy imaging [82, 199, 274]) such that the neural network can distinguish the foreign objects from the base objects. If changing the quality of the radiographs is not possible, a separate high-quality scan of the same object can be made under the same angles, to achieve more contrast of the foreign object in the reconstructions. The scanning routine can be carried out in any lab, as long as it is done under similar conditions as in the intended industrial X-ray imaging setting.

- **Reconstruction algorithm** (Fig. 2.3e): Depending on the type of data, different reconstruction algorithms may be considered [51, 112]. In this work, we have used the SIRT algorithm to account for the noise in the data, but other reconstruction algorithms such as Feldkamp-Davis-Kress (FDK) algorithm [87] or the Conjugate Gradient method for Least Squares (CGLS) [123] can be considered as well. Also, when dealing with spectral or generic multi-channel data, multi-channel reconstruction methods [143, 238, 246, 254, 312] can be used to increase the reconstruction accuracy even further. When dealing with objects that may change in time, dynamic reconstruction methods can be considered [71, 114, 211].
- **Segmentation algorithm** (Fig. 2.3f): In this work we have used a simple global thresholding scheme, but many more segmentation methods are available, as well as approaches to reduce possible noise [70], or bounding boxes when the location of the foreign object is known [144]. In case of multi-channel data, a multi-dimensional thresholding scheme can be used, as well as clustering methods. Discrete reconstructions algorithms that combine reconstruction and segmentation are also available [30, 121].
- **Virtual projection** (Fig. 2.3g): When creating the virtual projection, post-processing on the generated ground truth projections can be applied to increase the training target quality, for instance by denoising the obtained ground truth projections.
- **Supervised learning** (Fig. 2.3c and h): To validate the workflow, we have used the U-Net architecture with ADAM optimization on cross entropy loss and dice loss, as well as the MSD network with ADAM optimization [146] on the cross-entropy loss. Other neural network architectures (see Section 2.2.2) can also be considered, as well as different optimization strategies and loss functions. Note that the foreign object detection problem considered in this work may be ambiguous, since for a base object containing a foreign object another base object can theoretically be constructed (without foreign object) that results in the same radiograph. This constructed base object may have an unnatural shape when compared with other base objects, but if it happens, it may lead to inconsistent training data for the network. However, this possible problem is independent of the workflow and can be resolved by multi-spectral imaging or multi-angle imaging, and training the networks with multiple images from the same object resulting from these imaging

methods. However, creating reconstructions with data from these advanced imaging methods would not be necessary.

In Section 2.4, we have compared the segmentation results with workflow-generated data from many angles for each object to results with workflow-generated data with only one radiograph for each object. According to the results in Figures 2.8 and 2.14, segmentation and detection accuracy can be improved by using multiple annotated radiographs for each training object. For a comparative classical approach, ideally manual annotation of the data should be carried out, taking into account the variation that may occur in different annotations for the same set of radiographs. In addition to the tremendous effort required to manually annotate a large number of radiographs, there are a number of issues that arise in the food industry. To carry out segmentation on 2D radiographs, highly specific knowledge is required for which it is difficult to find experts, as opposed to trained radiologists in medical imaging. This is also reflected in the lack of suitable publicly available datasets with annotated radiographs of food products. Also, the manual annotations may vary depending on the detection goal by a manufacturer (such as object sizes, positions or number of foreign objects). For these reasons, we eventually chose not to compare the proposed data generation with manual annotation. Manual annotation could still be used in conjunction with this workflow to replace the segmentation step. Whether or not this is feasible and will yield better results will ultimately depend on the specific application at hand.

2.6 Conclusions

In this research, a new workflow is proposed for generating training data for supervised deep learning for foreign object detection in an industrial setting. In this workflow, a number of representative objects are scanned using X-ray imaging, reconstructed using computed tomography, segmented and virtually projected in an objective and reproducible manner to obtain the true foreign object locations in a large set of radiographs, after which supervised machine learning can be applied to detect foreign objects with high accuracy depending on the number representative objects included. We demonstrate this workflow on both laboratory and simulated data using neural networks for the deep learning task. Through laboratory experiments, we have verified that the workflow produces adequate target images. The introduced measures assess the quality of foreign object detection with networks trained using datasets generated with this workflow. All experiments show a consistent result in which the accuracy increases significantly with a few number of training objects, and less significantly for every additional training object. In the laboratory experiment, we consistently obtain high accuracies for detecting gravel in modelling clay with low exposure times using this workflow, demonstrating its application potential in an industrial setting.

Code and data availability

The source code to reproduce all results in this chapter are available on <https://github.com/mzeegers/DeepFODDataGenerator> [309]. These Python scripts make use of open-source toolboxes, in particular the ASTRA toolbox [1, 2], MSD-net [227] and PyTorch [222, 223]. The datasets generated for this chapter are available at Zenodo. Separate submissions are made for the processed data resulting in radiographs with ground truth for object detection [308], as well as the unprocessed CT scan data for complete reproduction of the results in this chapter [307].

