Multi-view X-ray R-CNN

Jan-Martin O. Steitz^[0000-0002-3549-312X], Faraz Saeedan^{\boxtimes [0000-0002-8932-2983]}, and Stefan Roth^[0000-0001-9002-9832]

Department of Computer Science, TU Darmstadt, Darmstadt, Germany faraz.saeedan@visinf.tu-darmstadt.de

Abstract. Motivated by the detection of prohibited objects in carry-on luggage as a part of avionic security screening, we develop a CNN-based object detection approach for multi-view X-ray image data. Our contributions are two-fold. First, we introduce a novel multi-view pooling layer to perform a 3D aggregation of 2D CNN-features extracted from each view. To that end, our pooling layer exploits the known geometry of the imaging system to ensure geometric consistency of the feature aggregation. Second, we introduce an end-to-end trainable multi-view detection pipeline based on Faster R-CNN, which derives the region proposals and performs the final classification in 3D using these aggregated multi-view features. Our approach shows significant accuracy gains compared to single-view detection while even being more efficient than performing single-view detection in each view.

1 Introduction

Baggage inspection using multi-view X-ray imaging machines is at the heart of most aviation security screening programs. Due to inherent shortcomings in human inspection arising from gradual fatigue, occasional erroneous judgments, and privacy concerns, computer-aided automatic detection of dangerous goods in baggage has long been a sought-after goal [19]. However, earlier approaches, mostly based on hand-engineered features and support vector machines, fell far short of providing detection accuracy comparable to human operators, which is critical due to the sensitive nature of the task [3, 8]. Thanks to recent advances in object detection using deep convolutional neural networks [10, 11, 22, 23] with stunning success in photographic images, the accuracy of single-view object detection in X-ray images has improved significantly [16]. Yet, most X-ray machines for baggage inspection provide multiple views (two or four views) of the screening tunnel. An example of this multi-view data is shown in Fig. 1. Multi-view approaches for these applications have been only used in the context of classical methods, but whether CNN-based detectors can benefit is unclear. In fact, [16] found that a naive approach feeding features extracted from multiple views simultaneously to a fully-connected layer for detection leads to a performance drop over the single-view case.

Fueled by applications such as autonomous driving, 3D object detection has lately gained momentum [9]. These 3D detection algorithms, unlike their 2D



Fig. 1. An example of multi-view X-ray images of hand luggage containing a glass bottle.

counterparts, are not general purpose and rely on certain sensor combinations or employ heavy prior assumptions that make them not directly applicable to multiview object detection in X-ray images. Some of these 3D detection algorithms assume that the shape of the desired object is known in the form of a 3D model [2,24]. Yet, 3D models of objects are more difficult to acquire compared to simple bounding box annotations and a detector that relies on them for training may not generalize well on objects with highly variable shape such as handguns. Other methods use point clouds from laser range finders alone or in conjunction with RGB data from a camera [5, 21]. Our setup, in contrast, provides multiview images of objects in a two channel (dual-energy) format, which is rather different from stereo or point cloud representations. Using 3D convolutions and directly extending existing 2D methods is one possibility, but the computational cost and memory requirements can be prohibitive when relying on very deep CNN backbone architectures such as ResNet [14].

In this work, we extend the well-known Faster R-CNN [23] to multi-view X-ray images by employing the idea of late fusion, which enables our use of very deep CNNs to extract powerful features from each individual view, while deriving region proposals and performing the classification in 3D. We introduce a novel *multi-view pooling layer* to enable this fusion of features from single views using the geometry of the imaging setup. This geometry is assumed fixed all through training and testing and needs to be calculated once. We do not assume further knowledge of the detected objects as long as we have sufficient bounding box annotations in 2D. We show that our method, termed MX-RCNN, is not only highly flexible in detecting various hazardous object categories with very little extra knowledge required, but is also considerably more accurate while even being more efficient than performing single-view detection in the four views.

2 Related Work

Inspired by the impressive results in 2D object detection, several recent works [20,30,32,33] build upon a 2D detection in an image and then attempt to estimate the 3D object pose. These methods, however, rely predominantly on a set of prior constraints such as objects being located on the ground for estimating the 3D position of the object. These prior constraints do not appear easily extensible to our problem case, since objects inside bags can take an arbitrary orientation and position.

Other approaches have tried to work directly with depth data [18, 27, 28], where most methods voxelize the space into a regular grid and apply 3D convolutions to the input. While this yields the most straightforward extension of well-proven 2D detectors, which are based on 2D convolution layers, increasing the dimensionality of convolution layers can only be done at very low resolution as denser voxelizations of the space result in unacceptably large memory consumption and processing time. Our method also uses the idea of 3D convolutions but defers them to very late stages in which we switch from 2D to 3D. In doing so, we enable the detector to leverage high image resolution in the input while extracting powerful features that leverage view-consistency.

A number of methods use the geometry of the objects as prior knowledge to infer a 6-DoF object pose. These methods mainly rely on CAD models or ground truth 3D object models and match either keypoints between these models and 2D images [2, 24, 34] or entire reconstructed objects [25]. Borrowing ideas from robotics mapping, the estimated pose of an object produced by a CNN can also be aligned to an existing 3D model using the iterative closest points (ICP) algorithm [13].

Hang *et al.* [29] proposed a method most closely related to ours, which aggregates features extracted from multiple views of a scene using an elementwise max operation among all views. Yet, unlike our multi-view pooling layer this aggregation is not geometry-aware.

Despite the recent focus on the problem of visual object detection, its application to X-ray images has not received as much attention. Some older methods [3, 8] exist for this application, but they perform considerably weaker than deep learning-based methods [1]. However, the use of CNNs on X-ray images for baggage inspection has been limited to the direct application of basic 2D detection algorithms, either with pretraining on photographic images or training from scratch. Jaccard *et al.* [16] propose a black-box approach to multi-view detection by extracting CNN features from all views, concatenating them, and feeding them to fully-connected layers. Yet, the accuracy fell short of that of the original single-view detection. To the best of our knowledge, there exists no previous end-to-end learning approach that successfully uses the geometry of the X-ray machine to perform a fusion of features from various views.

3 Multi-view X-ray R-CNN

We build on the standard Faster R-CNN object detection method [23], which works on single-view 2D images and is composed of two stages. They share a common feature extractor, which outputs a feature map. The first stage consists of a Region Proposal Network (RPN) that proposes regions of interest to the second stage. Those regions are then cut from the feature map and individually classified. The RPN uses a fixed set of 9 standard axis-aligned bounding boxes in 3 different aspect ratios and 3 scales, so-called anchor boxes. Those anchor boxes are defined at every location of the feature map. The RPN then alters their position and shape by learning regression parameters in training. Additionally,



Fig. 2. Schema of the hybrid 2D-3D MX-RCNN architecture. The features of each view are extracted independently in 2D and combined in the multi-view pooling layer (marked in red). The resulting common 3D feature volume is passed to the RPN and to the RoI pooling layer where regions are extracted for evaluation in the 3D R-CNN layers.

a score to distinguish between objects and background is learned in a classagnostic way, which is used for non-maximum suppression and to pass only a subset of top-scoring region proposals to the second stage. The second stage classifies the proposals and outputs additional regression values to fine-tune the bounding boxes at the end of the detection process.

MX-RCNN. The basic concept of our Multi-view X-ray R-CNN (MX-RCNN) approach is to perform feature extraction on 2D images to be able to utilize standard CNN backbones including ImageNet [26] pretraining. This addresses the fact that the amount of annotated X-ray data is significantly lower than that of photographic images. We then combine the extracted feature maps of different projections, or views, into a common 3D feature space in which object detection takes place, identifying 3D locations of the detected objects.

Our MX-RCNN uses a ResNet-50 [15] architecture, where the first 4 out of its 5 stages are used for feature extraction on the 2D images. Then a novel multi-view pooling layer, provided with the fixed geometry of the imaging setup, combines the feature maps into a common 3D feature volume.

The combined feature space is passed to a RPN, which has a structure similar to the RPN in Faster R-CNN [23], but with 3D convolutional layers instead of 2D ones. Further, it has $6 \times A$ regression parameter outputs per feature volume position, where A is the number of anchor boxes, because for 3D bounding boxes 6 regression parameters are needed. Following Faster R-CNN, we define those regression parameters as

$$\begin{aligned} t_x &= (x - x_a)/w_a & t_y &= (y - y_a)/h_a & t_z &= (z - z_a)/d_a \\ t_w &= \log(w/w_a) & t_h &= \log(h/h_a) & t_d &= \log(d/d_a) , \end{aligned}$$
(1)

where the index a denotes the parameters of the anchor box and with bounding box center (x, y, z), width w, height h, and depth d.

The RPN proposes volumes to be extracted by a Region-of-Interest pooling layer with an output of size $7 \times 7 \times 7$ to cover a part of the feature volume similar

in relative size to the 2D case. The 3D regions are then fed into a network similar to the last stage of ResNet-50 in which all convolutional and pooling layers are converted from 2D to 3D and the size of the last pooling kernel is adjusted to fit the feature volume size. In contrast to the 2D stages, these 3D convolutions are trained from scratch since ImageNet pretraining is not possible. Afterwards, the regions are classified and 3D bounding box regression parameters are determined. A schema of our MX-RCNN is depicted in Fig. 2.

3.1 K-means clustering of anchor boxes

When we expand the hand-selected aspect ratios of the Faster R-CNN anchor boxes of 1:1, 1:2, and 2:1 at 3 different scales to 3D, we arrive at a total of 21 anchor boxes. Since this number is large and limits the computational efficiency, we aim to improve over these standard anchor boxes. To that end, we assess the quality of the anchor boxes as priors for the RPN. Specifically, we use their intersection over union (IoU) [17] with the ground-truth annotations of the training set. We instantiate anchor boxes at each position of the feature map used by the RPN and for each ground-truth annotation we find its highest IoU with an anchor box. For the standard anchor boxes expanded to 3D, this yields an average IoU of 0.5.

To improve upon this while optimizing the computational efficiency, we follow the approach of the YOLO9000 object detector [22] and use k-means clustering on the bounding box dimensions (width, height, depth) in the training set to find priors with a better average overlap with the ground-truth annotations. We employ the Jaccard distance [17]

$$d_{\mathcal{J}}(a,b) = 1 - \operatorname{IoU}(a,b) \tag{2}$$

between boxes a and b as a distance metric. We run k-means clustering for various values of k. Fig. 3 shows the average IoU between ground-truth bounding boxes and the closest cluster (*blue, circles*). To convert clusters into anchor boxes, they have to be aligned to the resolution of a feature grid. To account for this, we also plot the IoU between the ground-truth bounding boxes and the closest cluster once it has been shifted to the nearest feature grid position (*red, diamonds*). We choose k = 10, which achieves an average IoU of 0.56 for the resulting anchor boxes distributed in a grid. This is clearly higher than using the standard 21 hand-selected anchor boxes, while maintaining the training and inference speed of a network with only 10 anchor boxes.

3.2 Multi-view pooling

Our proposed multi-view pooling layer maps the related feature maps of the 2D views of an X-ray recording into a common 3D feature volume by providing it with the known geometry of the X-ray image formation process in the form of a weight matrix. To determine the weights, we connect each group of detector locations related to one pixel in the 2D feature map to their X-ray source to form



Fig. 3. Average IoU between ground-truth boxes and the closest cluster from k-means clustering as a function of k. The IoU is calculated for anchor boxes centered on the ground-truth boxes (blue, circles) and for anchor boxes distributed in the grid of the feature map (*red, diamonds*). Even a low number of clusters already outperforms the standard hand-selected anchor boxes (IoU of 0.5).

beams across the 3D space. For each of the output cells of the 3D feature volume, we use the volume of their intersection with the beams normalized by the cell volume as relative weight factors. The multi-view pooling layer then computes the weighted average of the feature vectors of all beams for each output cell, normalized by the number of views in each X-ray recording; we call this variant MX-RCNN_{avg}. Additionally, we implemented a version of the multi-view pooling layer that takes the maximum across the weighted feature vectors of all beams for each output cell; we call this variant MX-RCNN_{max}. An example of a mapping, specific to our geometry, is shown in Fig. 4.

3.3 Conversion of IoU thresholds

Since the IoU in 3D (volume) behaves differently than in 2D (area), we aim to equalize for this. Specifically, we aim to apply the same strictness for spatial shifts that are allowed per bounding box dimension such that a proposed object is still considered a valid detection. We assume that the prediction errors of the bounding box regression values are equally distributed across all dimensions. For simplicity, we further assume that errors are only made up of shifts compared to the ground-truth bounding boxes. In 2D, the allowed relative shift s per dimension of a bounding box of arbitrary dimensions for an IoU threshold of t_2 is given by

$$s(t_2) = 1 - \sqrt{\frac{2t_2}{t_2+1}}$$
 (3)

The same relative shift applied to an arbitrarily sized 3D bounding box would require the threshold

$$t_3(s) = \frac{(1-s)^3}{(2-(1-s)^3)}.$$
(4)



Fig. 4. Example plot illustrating the relevant beams (*color-coded per view*) in the multi-view pooling of a specific output cell (*marked in red*). The actual geometry differs slightly.

An IoU threshold t_2 applied to the 2D case therefore becomes

$$t_3(t_2) = \frac{\sqrt{2t_2/(1+t_2)}^3}{2 - \sqrt{2t_2/(1+t_2)}^3} \tag{5}$$

for use with 3D bounding boxes if the same strictness per dimension is to be maintained. The evaluation of the PASCAL VOC challenge [7] for detection in 2D images uses a standard threshold of 0.5, which yields a threshold of 0.374 in the 3D detection case.

3.4 Computational cost

Our single-view Faster R-CNN implementation reaches a frame rate of 3.9 fps for training and 6.1 fps for inference on a NVIDIA GeForce GTX Titan X GPU. 4 frames or images need to be processed for one complete X-ray recording. Our MX-RCNN achieves frame rates of 4.6 fps for training and 7.6 fps for inference; note that it can share work across the 4 frames as they are processed simultaneously. Thus, our method is 18% faster in training and 25% faster in inference. We attribute this to the lower number of regions extracted per recording, because of the common classification stage in the multi-view detector, despite the higher computational costs for its 3D convolutional layers.



Fig. 5. False-color images of all views of an example X-ray recording of baggage containing a handgun. The handgun is easier to spot in certain images depending on their angle of projection. Bounding boxes show the original 2D annotations (black) and the reprojected 3D annotations (red).

4 Multi-view X-Ray Dataset

Lacking a standardized public dataset for this task, we leverage a custom dataset¹ of dual-energy X-ray recordings of hand luggage made by an X-ray scanner used for security checkpoints. The X-ray scanner uses line detectors located around the tunnel through which the baggage passes. Its pixels constitute the x-axis in the produced image data while the movement of the baggage, respectively the duration of the X-ray scan and the belt speed, define the y-axis of the image. Each recording consists of four different views, three from below and one from the side of the tunnel the baggage is moved through. The scans from each view produce two grayscale attenuation images from the dual-energy system, which are converted into a false-color RGB image that is used for the dataset. This is done to create images with 3 color channels that fit available models pretrained on ImageNet-1000 [26]. An example recording is shown in Fig. 5.

The following types of recordings are available in the dataset:

Glass Bottle. Recordings of baggage containing glass bottles of different shapes

¹ Unfortunately, we are not able to release the dataset to the public. Researchers wishing to evaluate on our dataset for comparison purposes are invited to contact the corresponding author.

Type / Subset	Train	Validation	Test	Total
Glass Bottle	358 (1432)	40 (160)	209 (836)	607 (2428)
TIP Weapon	$1944 \ (7776)$	216 (864)	0 (0)	2160 (8640)
Real Weapon	0 (0)	0 (0)	464 (1856)	464 (1856)
Negative	$0 \ (0)$	0 (0)	950 (3800)	950 (3800)
Total	2302 (9208)	256 (1024)	1623 (6492)	4181 (16724)

Table 1. Number of recordings (images) in the different subsets of the dataset.

and sizes.

TIP Weapon. Synthetic recordings of baggage where a pre-recorded scan of a handgun is randomly projected onto a baggage recording by a method called Threat Image Projection (TIP) [4]. A limited set of handguns is repeatedly used to generate all recordings.

Real Weapon. Recordings that contain a handgun of various types and are obtained using a conventional scan without the use of TIP.

Negative. Recordings containing neither handguns nor glass bottles.

The synthetic TIP images are only used for training and validation; the complete scans with real weapons are used for testing to evaluate if the trained network generalizes. A detailed overview of the different subsets of the data is given in Table 1. The dataset is split into its subsets such that views belonging to one and the same recording are not distributed over different subsets. PASCAL VOC-style annotations [7] with axis-aligned 2D bounding boxes per view are available for two different classes of hazardous objects, *weapon* and *glassbottle*.

4.1 3D bounding box annotations

To be able to train and evaluate our multi-view object detection with 3D annotations, we generate those out of more commonly available 2D bounding box annotations. Specifically, we generate axis-aligned 3D bounding boxes from several axis-aligned 2D bounding boxes. Because all our X-ray recordings have at most one annotated object,² there is no need to match multiple annotations across the different views. In case of multiple annotated objects per image, a geometrically consistent matching could be used.

Recall the specific imaging setup from above. If we now align the 2D y-axis (belt direction) to the 3D z-axis, the problem of identifying a suitable 3D bounding box reduces to a mapping from the 2D x-axis to a xy-plane in 3D. The lines of projection between the X-ray sources and the detectors corresponding to the x-axis limits of the 2D bounding boxes define the areas where the object could be located in the xy-plane per view. We intersect those triangular areas using the Vatti polygon clipping algorithm [31] and choose the minimum axis-aligned bounding box containing the resulting polygon as an estimation of the

² The number of annotated objects is a restriction of the dataset only; our detector is able to handle multiple objects per image.



Fig. 6. Example of the 3D bounding box generation. The lines of projection (color-coded per view) for the 2D bounding box positions overlap in the xy-plane of the 3D volume. The resulting polygon of their intersection (pink) is enclosed by the 3D bounding box (red). The actual geometry differs slightly.

object's position in 3D space. For the z-axis limits of the 3D bounding box, we take the mean of the y-limits of the 2D bounding boxes. An example of the generation process is shown in Fig. 6. Note that while our process of deriving 3D bounding boxes is customized to the baggage screening scenario, analogous procedures can be defined for more general imaging setups.

Note that in the ideal case, all projection lines would intersect in the 4 corners of the bounding box. However, due to variances in the annotations made independently for each view, this does not hold in practice. As a result, the bounding box enclosing the intersection polygon is an upper bound of the estimated position of the object in 3D space. We thus additionally project the 3D bounding boxes back to the 2D views to yield 2D bounding boxes that include the geometry approximation made in the generation of the 3D bounding box. The difference between an original 2D annotation and a reprojected 3D bounding box can be seen in Fig. 5. We use these 2D annotations to train a single-view Faster R-CNN as a baseline that assumes the same object localization as the multi-view networks trained on the 3D annotations. If more precise 3D bounding boxes are desired, they could be obtained from joint CT and X-ray recordings, which are becoming more common in modern screening machines.

5 Experiments

5.1 Training

We scale all images used for training and evaluation by a factor of 0.5 so they have widths of 384, 384, 352, and 416 px for the views 0 through 3 of each

X-ray recording. For all experiments we use a ResNet-50 [14] with parameters pretrained on ImageNet-1000 [26]. The first two stages of the ResNet-50 are fixed and we fine-tune the rest of the parameters. If not mentioned otherwise, hyper-parameters remain unchanged against the standard implementation of Faster R-CNN. All networks are trained by backpropagation using stochastic gradient descent (SGD).

As a baseline we train a standard implementation of single-view Faster R-CNN on the training set without data augmentation. Single-view Faster R-CNN detects objects on all views independently. We start with a learning rate of 0.001 for the first 12 epochs and continue with a rate of 0.0001 for another 3 epochs.

For training and evaluation of our MX-RCNN, we use a mini-batch size of 4 with all related views inside one mini-batch. We reduce the number of randomly sampled anchors in each 3D feature volume from 256 to 128 to better match the desired ratio of up to 1:1 between sampled positive and negative anchors. Additionally, we reduce the number of sampled RoIs to 64 to better match the desired fraction of 25% foreground regions that have an IoU of at least 0.5 with a 3D ground-truth bounding box. We accordingly reduce the learning rate by the same factor [12]. In the multi-task loss of the RPN, we change the balancing factor λ for the regression loss and the normalization by the number of anchor boxes to an empirically determined factor of $\lambda = 0.05$ without additional normalization. In practice, this ensures good convergence of the regression loss. We initialize all 3D convolutional layers as in the standard Faster R-CNN implementation. We then train MX-RCNN_{avg} on the training set for 28 epochs with a learning rate of 0.0005. A cut of the learning rate did not show any benefit when evaluating with the validation set. Additionally, we train MX-RCNN_{max} for 17 epochs with the same learning rate. Again, a learning rate cut did not show any further improvement on the validation set.

5.2 Evaluation criteria

Since there is no established evaluation criterion for our experimental setting. we are using the average precision (AP), which is the de-facto standard for the evaluation of object detection tasks. Specifically, we compare the trained networks with the evaluation procedure for object detection of the PASCAL VOC challenge that is in use since 2010 [6]. To be considered a valid detection, a proposed bounding box must exceed a certain threshold with a ground-truth bounding box. If multiple proposed bounding boxes match the same groundtruth bounding box, only the one with the highest confidence is counted as a valid detection. The PASCAL VOC challenge uses an IoU of 0.5 as threshold for the case of object detection in 2D images. As discussed in Section 3.3, to apply the same strictness for relative shifts that are allowed per dimension, we set this threshold to 0.374 for the evaluation of 3D bounding boxes per recording. Nevertheless, the corresponding IoU threshold we derived for 3D is an estimate. Hence, we additionally project the proposed 3D bounding boxes onto the 2D views and evaluate them per image with a threshold of 0.5 to directly compare to standard 2D detection.

12 J.-M. O. Steitz et al.

Table 2. Experimental results of the different networks evaluated on the test set. For the multi-view networks the evaluation was done with the proposed 3D bounding boxes and their projections onto the 2D views.

Method	Single-view	$\mathrm{MX} ext{-}\mathrm{RCNN}_\mathrm{avg}$		$\mathrm{MX} ext{-}\mathrm{RCNN}_{\mathrm{max}}$	
Evaluation	2D	3D	2D	3D	2D
Weapon AP	85.56~%	92.28~%	90.32~%	89.01~%	87.73 %
Glassbottle AP	96.90 %	98.84 %	95.37~%	98.74~%	95.62~%
Mean AP	91.23~%	95.56~%	92.84~%	93.88~%	91.68~%



Fig. 7. Precision-recall curves of the different networks evaluated on the test set. The plots show the precision-recall curves of the single-view (*orange*, *dashed*), MX-RCNN_{avg} (*turquoise*, *solid*) and MX-RCNN_{max} (*pink*, *dotted*) networks. For the multi-view networks the precision-recall curves are shown for the evaluation in 3D.

5.3 Experimental results

The standard single-view Faster R-CNN [23] used as a baseline reaches 91.2% mean average precision (mAP) per image on the test set (average over classes). With 93.9% mAP per recording, our MX-RCNN_{max} is 2.7% points better than the baseline when evaluating in 3D and 0.5% points better when projected to 2D (91.7% mAP per image). Using a weighted average in the multi-view pooling layer of MX-RCNN (MX-RCNN_{avg}) shows consistently better detection accuracy than MX-RCNN_{max} except for the AP of the glassbottle class projected to 2D. With an mAP per recording for 3D evaluation of 95.6%, our MX-RCNN_{avg} is 4.4% points better than the baseline when evaluating in 3D and 1.7% better when projected to 2D (mAP of 92.9%). Detailed numbers of the evaluation can be found in Table 2. Additionally, precision-recall curves are provided in Fig. 7 to allow studying the accuracy at different networks within each object class is consistent with their AP values. Moreover, it becomes apparent that the benefit

Table 3. Tolerance of MX-RCNN for disabling individual views when evaluating on the test set. The proposed 3D bounding boxes were directly evaluated without reprojection to 2D.

Views	all	w/o bottom left	w/o bottom right	w/o right side			
Weapon AP	92.28~%	82.34 %	85.57~%	53.19~%			
Glassbottle AP	98.84~%	91.58~%	98.59~%	92.42~%			
Mean AP	95.56~%	86.96~%	92.08~%	72.80~%			
(b) MX-RCNN _{max}							
Views	all	w/o bottom left	w/o bottom right	w/o right side			
Weener AD	20.01.07	74.40.07	02 C0 07	F0.0C 07			
weapon AP	89.01 %	74.40 %	83.09 %	58.80 %			
Glassbottle AP	98.74~%	89.89~%	97.93~%	77.26~%			
Mean AP	93.88~%	82.14~%	90.81~%	68.06~%			

(a) MX-RCNN_{avg}

of the proposed multi-view approach is particularly pronounced in the important high-recall regime.

5.4 Ablation study

To test the importance of the different views in our multi-view setup on the final detection result, we disabled individual views while evaluating and in the mapping provided to the multi-view pooling layer, respectively. The evaluation was done for both variants of the multi-view pooling layer on the test set and we directly compared the proposed 3D bounding boxes to 3D ground-truth annotations with 0.374 as IoU threshold without reprojecting them to 2D. The detailed results can be found in Table 3. We notice that the tolerance for missing views from below the tunnel is higher than for a missing side view. Also, the impact is more pronounced on weapons whose appearance is more affected by out-of-plane rotations. In general, the use of weighted averaging when combining features in the multi-view pooling seems to be more fault tolerant than the use of a weighted maximum, with the exception of weapons in combination with a missing side view. The results show that the network indeed relies on all views to construct the feature volume and to propose and validate detections.

6 Conclusion

In this paper we have introduced MX-RCNN, a multi-view end-to-end trainable object detection pipeline for X-ray images. MX-RCNN is a two stage detector similar to Faster R-CNN [23] that extracts features from all views separately using a standard CNN backbone and then fuses these features together to shape 14 J.-M. O. Steitz et al.

a 3D representation of the object in space. This fusion happens in a novel multiview pooling layer, which combines all individual features leveraging the geometry of the X-ray imaging setup. An experimental analysis on a dataset of carry-on luggage containing glass bottles and hand guns showed that when trained with the same annotations, MX-RCNN outperforms Faster R-CNN applied to each view separately and is computationally cheaper than separate processing of all views. We also showed in an ablation study that the method works by far better when the view angles do not all fall in one line (degenerate 3D case), showing that the pipeline is indeed leveraging its 3D feature representation.

Acknowledgements. The authors gratefully acknowledge support by Smiths Heimann GmbH.

References

- Akçay, S., Kundegorski, M.E., Devereux, M., Breckon, T.P.: Transfer learning using convolutional neural networks for object classification within X-ray baggage security imagery. In: ICIP. pp. 1057–1061 (2016). https://doi.org/10.1109/ICIP.2016.7532519
- Aubry, M., Maturana, D., Efros, A.A., Russell, B., Sivic, J.: Seeing 3D chairs: Exemplar part-based 2D-3D alignment using a large dataset of CAD models. In: CVPR. pp. 3762–3769 (2014). https://doi.org/10.1109/CVPR.2014.487
- Baştan, M.: Multi-view object detection in dual-energy X-ray images. Mach. Vision Appl. 26(7-8), 1045–1060 (Nov 2015). https://doi.org/10.1007/s00138-015-0706-x
- Brudy, T., Schilhab, S.: Projection of hazardous items into X-ray images of inspection objects. Patent WO 2016/001282 AI (Jan 2016)
- Chen, X., Ma, H., Wan, J., Li, B., Xia, T.: Multi-view 3D object detection network for autonomous driving. In: CVPR. pp. 1907–1915 (2017). https://doi.org/10.1109/CVPR.2017.691
- Everingham, M., Eslami, S.M.A., Van Gool, L., Williams, C.K.I., Winn, J., Zisserman, A.: The PASCAL visual object classes challenge: A retrospective. Int. J. Comput. Vision 111(1), 98–136 (2014). https://doi.org/10.1007/s11263-014-0733-5
- Everingham, M., Van Gool, L., Williams, C.K.I., Winn, J., Zisserman, A.: The PASCAL visual object classes (VOC) challenge. Int. J. Comput. Vision 88(2), 303–338 (2010). https://doi.org/10.1007/s11263-009-0275-4
- Franzel, T., Schmidt, U., Roth, S.: Object detection in multi-view X-ray images. In: Pinz, A., Pock, T., Bischof, H., Leberl, F. (eds.) DAGM/OAGM 2012, LNCS, vol. 7476, pp. 144–154. Springer, Berlin, Heidelberg (2012). https://doi.org/10.1007/978-3-642-32717-9_15
- Geiger, A., Lenz, P., Urtasun, R.: Are we ready for autonomous driving? The KITTI vision benchmark suite. In: CVPR. pp. 3354–3361 (2012). https://doi.org/10.1109/CVPR.2012.6248074
- 10. Girshick, R.: Fast R-CNN. In: ICCV. pp. 1440–1448. IEEE (2015). https://doi.org/10.1109/ICCV.2015.169
- Girshick, R., Donahue, J., Darrell, T., Malik, J.: Region-Based convolutional networks for accurate object detection and segmentation. IEEE T. Pattern Anal. Mach. Intell. 38(1), 142–158 (Jan 2016). https://doi.org/10.1109/TPAMI.2015.2437384

- Goyal, P., Dollár, P., Girshick, R., Noordhuis, P., Wesolowski, L., Kyrola, A., Tulloch, A., Jia, Y., He, K.: Accurate, large minibatch SGD: Training ImageNet in 1 hour. arXiv:1706.02677 (2017)
- Gupta, S., Arbeláez, P., Girshick, R., Malik, J.: Aligning 3D models to RGB-D images of cluttered scenes. In: CVPR. pp. 4731–4740 (2015). https://doi.org/10.1109/CVPR.2015.7299105
- He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: CVPR. pp. 770–778 (2016). https://doi.org/10.1109/CVPR.2016.90
- He, K., Zhang, X., Ren, S., Sun, J.: Identity mappings in deep residual networks. In: Leibe, B., Matas, J., Sebe, N., Welling, M. (eds.) ECCV 2016, LNCS, vol. 9908, pp. 630–645. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-46493-0_38
- Jaccard, N., Rogers, T.W., Morton, E.J., Griffin, L.D.: Automated detection of smuggled high-risk security threats using deep learning. ICDP 2016 pp. 1–6 (Nov 2016). https://doi.org/10.1049/ic.2016.0079
- Jaccard, P.: The distribution of the flora in the Alpine zone. New Phytol. 11(2), 37–50 (Feb 1912). https://doi.org/10.1111/j.1469-8137.1912.tb05611.x
- Li, B.: 3D fully convolutional network for vehicle detection in point cloud. In: IROS. pp. 1513–1518 (2016). https://doi.org/10.1109/IROS.2017.8205955
- Mery, D., Svec, E., Arias, M., Riffo, V., M. Saavedra, J., Banerjee, S.: Modern computer vision techniques for X-ray testing in baggage inspection. IEEE Transactions on Systems, Man, and Cybernetics: Systems 15(2), 1–11 (Dec 2016). https://doi.org/10.1109/TSMC.2016.2628381
- Mousavian, A., Anguelov, D., Flynn, J., Košecká, J.: 3D bounding box estimation using deep learning and geometry. In: CVPR. pp. 5632–5640 (2017). https://doi.org/10.1109/CVPR.2017.597
- Qi, C.R., Liu, W., Wu, C., Su, H., Guibas, L.J.: Frustum PointNets for 3D object detection from RGB-D data. In: CVPR. pp. 918–927 (2018)
- Redmon, J., Farhadi, A.: YOLO9000: Better, faster, stronger. CVPR (2017). https://doi.org/10.1109/CVPR.2017.690
- Ren, S., He, K., Girshick, R., Sun, J.: Faster R-CNN: Towards real-time object detection with region proposal networks. IEEE T. Pattern Anal. Mach. Intell. 39(6), 1137–1149 (Jun 2017). https://doi.org/10.1109/TPAMI.2016.2577031
- Romea, A.C., Torres, M.M., Srinivasa, S.: The MOPED framework: Object recognition and pose estimation for manipulation. Int. J. Robot. Res. 30(10), 1284–1306 (Sep 2011). https://doi.org/10.1177/0278364911401765
- Rothganger, F., Lazebnik, S., Schmid, C., Ponce, J.: 3D object modeling and recognition using local affine-invariant image descriptors and multiview spatial constraints. Int. J. Comput. Vision 66(3), 231–259 (2006). https://doi.org/10.1007/s11263-005-3674-1
- Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., Berg, A.C., Fei-Fei, L.: ImageNet large scale visual recognition challenge. Int. J. Comput. Vision 115(3), 211–252 (2015). https://doi.org/10.1007/s11263-015-0816-y
- Song, S., Xiao, J.: Sliding shapes for 3D object detection in depth images. In: ECCV. pp. 634–651 (2014). https://doi.org/10.1007/978-3-319-10599-4_41
- Song, S., Xiao, J.: Deep sliding shapes for amodal 3D object detection in RGB-D images. In: CVPR. pp. 808–816 (2016). https://doi.org/10.1109/CVPR.2016.94
- Su, H., Maji, S., Kalogerakis, E., Learned-Miller, E.G.: Multi-view convolutional neural networks for 3D shape recognition. In: ICCV. pp. 945–953 (2015). https://doi.org/10.1109/ICCV.2015.114

- 16 J.-M. O. Steitz *et al.*
- 30. Tulsiani, S., Malik, J.: Viewpoints and keypoints. In: CVPR. pp. 1510-1519 (2015)
- Vatti, B.R.: A generic solution to polygon clipping. Commun. ACM 35(7), 56–63 (Jul 1992). https://doi.org/10.1145/129902.129906
- 32. Xiang, Y., Choi, W., Lin, Y., Savarese, S.: Data-driven 3D voxel patterns for object category recognition. In: CVPR. pp. 1903–1911 (2015). https://doi.org/10.1109/CVPR.2015.7298800
- Xiang, Y., Choi, W., Lin, Y., Savarese, S.: Subcategory-aware convolutional neural networks for object proposals and detection. In: WACV. pp. 924–933 (2017). https://doi.org/10.1109/WACV.2017.108
- 34. Zhu, M., Derpanis, K.G., Yang, Y., Brahmbhatt, S., Zhang, M.M., Phillips, C.J., Lecce, M., Daniilidis, K.: Single image 3D object detection and pose estimation for grasping. In: ICRA. pp. 3936–3943 (2014). https://doi.org/10.1109/ICRA.2014.6907430