

Operationalizing Convolutional Neural Network Architectures for Prohibited Object Detection in X-Ray Imagery

Thomas W. Webb¹, Neelanjan Bhowmik¹, Yona Falinie A. Gaus¹, Toby P. Breckon^{1,2}
Department of {Computer Science¹ | Engineering²}, Durham University, UK

Abstract—The recent advancement in deep Convolutional Neural Network (CNN) has brought insight into the automation of X-ray security screening for aviation security and beyond. Here, we explore the viability of two recent end-to-end object detection CNN architectures, Cascade R-CNN and FreeAnchor, for prohibited item detection by balancing processing time and the impact of image data compression from an operational viewpoint. Overall, we achieve maximal detection performance using a FreeAnchor architecture with a ResNet₅₀ backbone, obtaining mean Average Precision (mAP) of 87.7 and 85.8 for using the OPIXray and SIXray benchmark datasets, showing superior performance over prior work on both. With fewer parameters and less training time, FreeAnchor achieves the highest detection inference speed of ~ 13 fps (3.9 ms per image). Furthermore, we evaluate the impact of lossy image compression upon detector performance. The CNN models display substantial resilience to the lossy compression, resulting in only a 1.1% decrease in mAP at the JPEG compression level of 50. Additionally, a thorough evaluation of data augmentation techniques is provided, including adaptations of MixUp and CutMix strategy as well as other standard transformations, further improving the detection accuracy.

Index Terms—X-ray imagery, deep convolutional neural networks, object detection, lossy compression, data augmentation.

I. INTRODUCTION

X-ray security screening is an extensively employed method within aviation and broader transportation. Prohibited item detection via scanned X-ray imagery provides a non-intrusive internal view of scanned baggage, freight and postal items. A modern X-ray security scanner [1] makes use of dual X-ray energy to acquire images at two discrete energy levels and subsequently combine the information with a help of colour transformation function to produce a single pseudo-colour X-ray image (Figure 1). Undertaking manual detection using the pseudo-colour X-ray imagery presents a significant image-based screening task for human operators. The cluttered nature of the X-ray imagery or the inclusion of compact, complex structures can seriously impact human operator performance. For instance, a prohibited item which is concealed within an electronics item can become very challenging to recognise, due to the compact structure of a laptop, limiting human detection capabilities [2]. Human operator performance can be further impacted by external factors such as emotional exhaustion or job satisfaction that adversely affect manual screening accuracy [3]. From examining expert human detec-

tion performance, the work of [3] suggests that an accuracy of only 80 – 90% is achieved by humans.

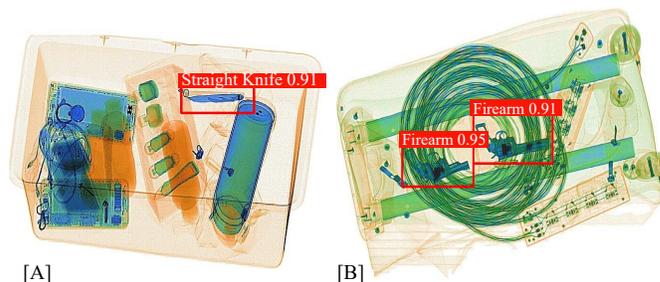


Fig. 1. Exemplar X-ray imagery from [A] OPIXray [4] and [B] SIXray [5] containing prohibited item detection examples in bounding boxes.

Prior work on prohibited item detection in X-ray security imagery originally uses traditional computer vision methods that rely on handcrafted features such as Bag of Visual Words (BoVW) [6], [7]. The work of [6] investigates the performance of a range of feature point descriptors within a classical BoVW model, supported by SVM and Random Forest classification, achieving 94% accuracy on a two-class firearm detection problem with SVM classification. Subsequent advancements of deep learning have brought insight into the automation of this X-ray screening task, with Convolutional Neural Network (CNN) based methods proven to be effective in detecting a wide range of object classes within X-ray screening [4], [5], [8]–[10]. The work of [8] extends to examine object detection, evaluating the performance of a sliding window CNN detector against contemporary region-based detection strategies, Faster R-CNN [11], R-FCN [12] and YOLOv2 [13]. These contemporary models outperform the traditional sliding window approach, achieving a maximal 0.88 and 0.97 mean Average Precision (mAP) over a six-class object detection and two-class firearm detection problem respectively. The work of [10] addresses the limited availability of public datasets of X-ray security imagery by exploring the impact of using synthetic X-ray training imagery for CNN architectures. Faster R-CNN [11] and RetinaNet [14] are trained on real and synthetic X-ray training imagery with Faster R-CNN achieving a maximal performance of 0.88 mAP when trained and tested with real X-ray training data on a three-class detection problem.

CNN based methods have proven effective in detecting prohibited items in X-ray security imagery [8], [9], [15]. However, the performance of such object detection approaches poses three main issues. Firstly, it is heavily reliant on the

availability of a substantial volume of labelled X-ray imagery. The work of [5] introduces a large-scale dataset, *SIXray*, fully illustrating the extremely cluttered nature of X-ray security threats imagery (Figure 1, B). Another public dataset, *OPIXray* [4] focuses on the widely occurring prohibited items (Figure 1, A). The availability of such X-ray datasets suitable for training CNN architectures is limited and also restricted in size and item signature [10].

The second considers the practical implementation of CNN architectures at security screening checkpoints. Such an implementation could be achieved through remote screening but the quantity of data involved during training and live application would require substantial storage and transmission infrastructure. Lossy compression techniques are used to reduce such overheads but lead to image quality degradation that can affect the performance of CNN [16]. The work of [17] evaluates the effect of JPEG compression, among other quality distortions, using CNN architectures when classifying a subset of the ImageNet dataset [18]. The work of [16] thoroughly evaluates the impact of JPEG and H.264 lossy compression have on CNN architectures. They evaluate Faster R-CNN [11] upon the Pascal VOC dataset [19] and find similarly to [17] that performance degrades rapidly at high lossy compression levels.

Thirdly, the use of CNN-based architectures result in the requirement for large datasets to avoid overfitting, but collating such large volumes of X-ray security imagery remains a significant challenge. This makes data augmentation, describing the set of techniques that can be used to increase the size, quality and variation of the training data, a particularly powerful tool within this domain. This is proven to reliably improve network generalisation [20] but selecting appropriate techniques requires prior knowledge of the dataset. There has been substantial investigation of data augmentation techniques within photographic imagery [21]–[23] but limited work has been done examining the performance of these techniques applied to X-ray imagery.

Addressing the above issues, in this work we make the following contributions:

- We investigate two end-to-end CNN object detection architectures, which differ in design, i.e., multi-stage (Cascade R-CNN [24]) vs one-stage (FreeAnchor [25]), for prohibited item detection by balancing processing time and performance from an operational viewpoint.
- As dual energy pseudo-colour X-ray imagery exhibits inherently different properties from standard photographic imagery, we examine the impact of using lossy image compression [26] on CNN based prohibited item detection, this may form a key component of a remote screening solution.
- We evaluate a range of data augmentation techniques, including adaptations of the MixUp [23] and CutMix [22] strategies, as a method for improving CNN-based object detector performance for prohibited item detection.

II. PROPOSED APPROACH

A. Detection Strategies

We consider two region based CNN detectors, Cascade R-CNN [24] and FreeAnchor [25] based on their high performance on public object detection benchmarks [27].

Cascade R-CNN (CR-CNN) [24], is an expansion of the Faster R-CNN [11]. It is a multi-stage architecture, where the output of each stage is fed into the next one for higher quality refinement (Figure 2, A). Due to the difficulty in producing a single regressor to perform optimally over all Intersection over Union (IoU) thresholds, the training data of each stage is sampled with increasing IoU thresholds, which inherently handles different training distributions. This enables the cascaded regression to be used as a re-sampling technique to incrementally improve the IoU of bounding boxes produced at each stage, containing a classifier and regressor, with each regressor accepting the bounding boxes produced from the previous stage.

FreeAnchor (FA) [25], adapts anchor-based detectors from hand-crafted assignment of anchors to a learned anchor-matching approach. It is a one-stage detector, which proposes anchor-object matching as a maximum likelihood estimation procedure, selecting the most representative anchor from a bag of candidate anchors generated for each object (Figure 2, B). Unlike previous CNN driven object detectors where the anchor is assigned for ground-truth objects under the restriction of object-anchor IoU, in this approach it proposed a learning-to-match approach and breaks the IoU restriction, which allows objects to match anchors in a flexible manner.

B. Data Augmentation

We examine a large collection of data augmentation techniques, in order to investigate their impact on X-ray prohibited item detection. We first consider a set of data-agnostic image transformations applied to the uncompressed datasets. We examine three sets of augmentations: geometric transformations, global image processing transformations, and methods involving combining images to generate new samples (Figure 3). For geometric transformations we select: random vertical and horizontal flipping (RandomFlip), a random crop of the image (RandomCrop), and rotating the image (Rotate). For global image processing transformations we select: blurring (Blur), equalising the image histogram (Equalise), and applying JPEG compression (JPEG). For the final set of augmentations we select MixUp [23] and CutMix [22] based on their performance when applied to photographic image classification and object detection tasks respectively. We adapt and implement these techniques for X-ray prohibited item detection.

MixUp. For our variation of MixUp [23] after the input image (x_i) is sampled from a dataset, a second image (y_i) is sampled through the same method. Both images are resized to identical dimensions. A new image (\hat{x}) is generated from the linear interpolation of the image pair (Eq. 1), controlled by the factor of $\lambda \in [0, 1]$. For natural image classification the one-hot encoded class vectors would be interpolated in the fashion

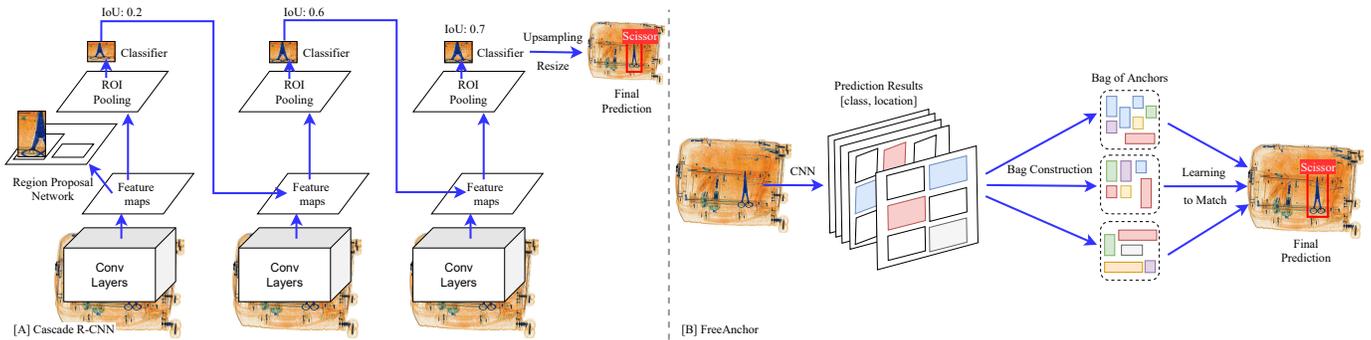


Fig. 2. Architecture of the CNN based detectors evaluated: [A] Cascade R-CNN [24], [B] FreeAnchor [25].

to produce the generated image label. For the X-ray domain we argue that this produces similar images to real samples and instead we combine the class vectors of both images as well as using the bounding boxes of both images as the targets. We use $\lambda = 0.5$ since both samples are equally valuable during training when specific classes are not being targeted.

$$\hat{x} = \lambda x_i + (1 - \lambda)x_j \quad (1)$$

We further adapt MixUp to produce a variation that can target specific classes at a bounding box level (*BboxMixUp*). Rather than interpolating between two entire images this variation selects one object from the second image x_j and interpolates between the bounding box area of the object in x_j and the respective area in x_i , after which the area in x_i is replaced with the newly interpolated sub-section. This generates a sample with a consistent background complexity while still introducing a new object.

CutMix. CutMix [22] has improved the performance of object detectors but it has only been designed for use in the CNN backbone pre-training step. CutMix takes two samples and replaces one rectangular area in one image with the respective area in the second image to generate a new samples. The target classification vector is linearly interpolated based on the proportion from each image. In order to apply CutMix directly during object detection training we instead select an object from each image and apply the same methodology the bounding boxes of those objects. We select an object bounding box b_i from x_i and another object bounding box b_j from x_j , resizing b_j to the same dimensions as b_i . We then replace a proportion b_i with the respective area of b_j using Eq. 2 to produce \hat{b} and insert this area back into x_i to generate the new sample \hat{x} . $\mathbf{M} \in \{0, 1\}$ is a binary mask describing the areas to remove and combine for the two bounding boxes. We select \mathbf{M} to denote a proportion of 0.5 of each bounding box to maximise the likelihood that there are sufficient features present for both objects.

$$\hat{b} = \mathbf{M} \odot b_i + (1 - \mathbf{M}) \odot b_j \quad (2)$$

The object from x_j for MixUp and the objects from each image in CutMix are selected based upon having IoU values with other objects in their original image below a defined threshold. In this work we use a threshold of 0.3. This is to avoid introducing partial features to the generated sample

without including the respective annotation targets, this could degrade model quality during training. The particular class to introduce in MixUp or the class pairings to combine in CutMix can be specified allowing training to target poorly performing classes for a given detector. We include results for this class specified variation of CutMix (*ClassCutMix*).

III. EXPERIMENTAL SETUP

A. Datasets

OPIXray [4]. consists of 8,885 dual-energy X-ray images with five-classes (*folding knife, straight knife, scissor, utility knife, multi-tool knife*), and images mimic the randomly stacked and overlapping reality of personal baggage.

SIXray [5]. consists of dual-energy X-ray images with six classes containing prohibited items. The images are of baggage and parcel X-ray scans collected from several subway stations where the distribution of the general items included corresponds to stream-of-commerce occurrence. We use a subset of this dataset, which is composed of five-classes (*gun, knife, wrench, pliers, scissors*) and total of 8,929 images.

Compressed datasets. To determine how much lossy JPEG compression [26] is viable within CNN object detection architecture in the X-ray security domain, both uncompressed *SIXray* [5] and *OPIXray* [4] datasets are compressed at three levels $\{95, 50, 10\}$ to create compressed version of the datasets (*OPIXray_c* and *SIXray_c*). Initially, the models are trained on uncompressed imagery and evaluated on the compressed imagery. Each architecture is subsequently retrained with compressed imagery at each of the three levels of lossy compression to determine whether resilience to compression could be improved, and how much compression we can afford before a significant impact on performance is observed.

B. Implementation detail

The CNN architectures (Section II-A) are implemented using MMDetection framework [28]. All experiments are initialised with weights pretrained on ImageNet [18]. Our CNN architectures (Section II-A) are trained using ResNet₅₀ [29] backbone on a Nvidia Titan XP GPU with following training configuration: backpropagation optimisation performed via Stochastic Gradient Descent (SGD), with initial learning rates of $2.5e - 4$ and $3.75e - 5$ respectively both with decay by a factor of 10 at 7th epoch and a batch size of 4.

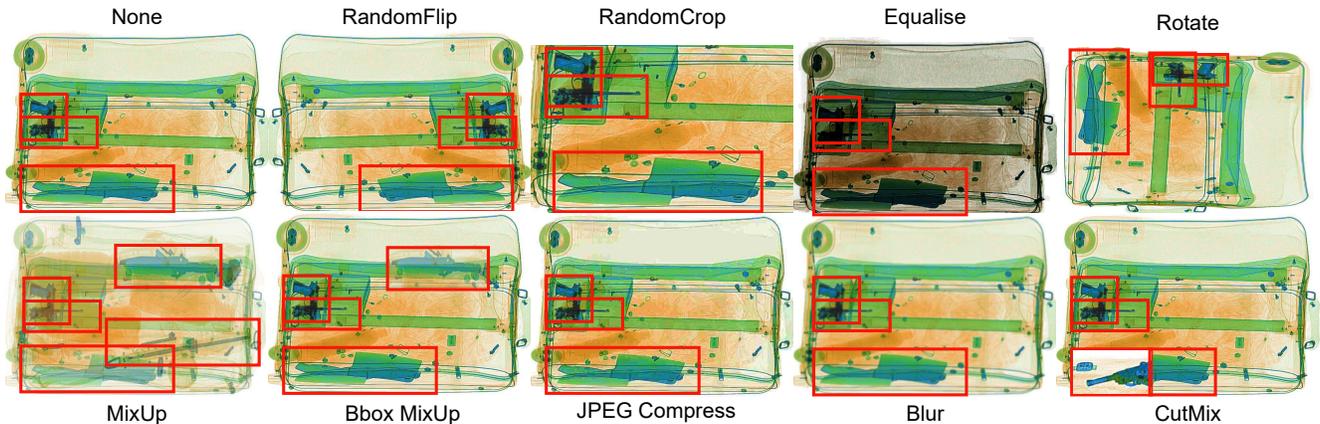


Fig. 3. Example of data augmentation techniques applied to X-ray security images.

IV. EVALUATION

The model performance is evaluated by MS-COCO metrics [27] (IoU > 0.50), using Average Precision (AP) for class-wise and mAP for the overall performance. Additionally, we use Complexity (number of parameters in millions, C), the ratio between mAP and the number of parameters in the architecture (mAP:C), model training time (in hours, T) and achievable frames per second (fps) throughput for comparing the model performance. The highlighted values in each table denote the maximal performance achieved.

A. Benchmarking on X-ray dataset

TABLE I
PERFORMANCE (CLASS NAME INDICATES AP, MAP INDICATES MEAN AVERAGE PRECISION OF ALL CLASSES) OF CNN MODELS.

Model	Folding	Straight	Scissor	Utility	M-tool	mAP
<i>OPIXray</i> CR-CNN [24]	90.6	70.5	97.0	79.7	91.4	85.9
FA [25]	91.7	74.6	97.3	80.8	94.1	87.7
FCOS+DOAM [4]	86.7	68.6	90.2	78.8	87.7	82.4
Model	Firearm	Knife	Wrench	Pliers	Scissors	mAP
<i>SIXray</i> CR-CNN [24]	87.4	80.4	82.4	86.3	86.4	84.6
FA [25]	89.8	81.0	79.6	88.6	89.9	85.8
CHR [5]	87.1	88.3	85.9	66.1	70.5	79.6

Table I presents prohibited items detection results for the first set of experiments using the CNN architectures set out in Section II are applied to *OPIXray* [4] and *SIXray* [5] to provide benchmark performances. We observe that the best performance for *OPIXray* (mAP: 87.7, Table I, upper) is achieved by FreeAnchor [25] producing maximal AP across all five classes. Cascade R-CNN [24] achieves a comparable mAP of 85.9 and both models significantly outperform the prior works FCOS+DOAM [4] (mAP of 82.4). For *SIXray*, FreeAnchor [25] again produces the best performance outperforming Cascade R-CNN [24] (mAP: 85.8 vs 84.6, Table I, lower). Both models again outperform the prior work of CHR [5] in terms of mAP. Overall, FreeAnchor [25] achieves the maximal mAP for both datasets. This is possibly because FreeAnchor [25] does not use fixed hand-crafted anchors during detection, and it is not practical to match the anchors using predefined sizes when the prohibited items are often slender objects (*knives*, *scissors*), cluttered and overlapped under X-ray scanner. Instead, it selects the most suitable anchor from

a set of candidate anchors by employing likelihood estimation for each target prohibited item.

TABLE II

Dataset	MODEL COMPLEXITY AND EFFICIENCY.				
	Model	C	mAP:C	T	fps
<i>OPIXray</i>	CR-CNN [24]	69.2	1.24	10.9	9.5
	FA [25]	36.4	2.41	7.8	10.0
<i>SIXray</i>	CR-CNN [24]	69.2	1.22	7.3	12.5
	FA [25]	36.4	2.36	7.1	13.3

Along with AP/mAP, Table II presents the computational efficiency, training time (T, hours taken to train for 30 epochs) and speed, which are crucial criteria for operational perspective and real-world deployment applications. FreeAnchor [25] has $\sim 2\times$ fewer parameters than Cascade R-CNN [24] while always achieving a higher computational efficiency (mAP:C) of 2.41 and 2.36 on *OPIXray* [4] and *SIXray* [5], respectively. Both models achieve real-time operational throughput (fps), which is within the belt speed (0.2 meter/second) of a standard X-ray scanner [1]. FreeAnchor [25] achieves $\sim 1.08\times$ the throughput of Cascade R-CNN [24] on each dataset (Table II). Demonstrating the trade-off that FreeAnchor [25] has made to exceed multi-stage Cascade R-CNN [24] performance. While comparing the training time, FreeAnchor [25] requires the shortest time (hours) to train for both *OPIXray* and *SIXray* (T: 7.8 and T: 7.1, respectively), and trains on average in 80% of the time required to train Cascade R-CNN [24] while achieving a higher overall detection accuracy.

B. Effect of Lossy Compression

Table III presents the results from evaluating CNN models with the varying optimiser and using two X-ray datasets after undergoing lossy compression (JPEG) for quality levels of: {10, 50, 95}. The models are first trained on the uncompressed datasets (*OPIXray*, *SIXray*) and evaluated on the compressed variants (*OPIXray_c*, *SIXray_c*). Subsequently, the models are re-trained with the compressed imagery at the respective level and evaluated on the compressed variants. When trained on the uncompressed dataset, the models retain similar performances at 50% and 95% compression levels with mAP only decreasing by 0.2% and 0.6% on average for *OPIXray*, and *SIXray* respectively. A compression level of 10% has a greater impact with mAP decreasing by 10.3% and 7.9% on the respective dataset. Retraining the models on imagery lossily compressed at the

same level (*OPIXray_c*, *SIXray_c*) significantly ameliorated the performance for *OPIXray_c* (Table III), but has a minimal effect for models applied to *SIXray_c*. The benefit of retraining the models on compressed images is illustrated in Figure 4, B, where compressed image trained models successfully detect the prohibited item, contrary to the missing detection (in white box) in Figure 4, A.

TABLE III
EFFECT OF LOSSY COMPRESSION ON DETECTION PERFORMANCE (mAP).

Train \Rightarrow Evaluation	Model	Lossy compression		
		10	50	95
<i>OPIXray</i> \Rightarrow	CR-CNN [24]	78.2	85.6	85.9
<i>OPIXray_c</i>	FA [25]	77.4	87.2	87.4
<i>OPIXray_c</i> \Rightarrow	CR-CNN [24]	84.9	86.1	85.2
<i>OPIXray_c</i>	FA [25]	87.4	87.6	88.1
<i>SIXray</i> \Rightarrow	CR-CNN [24]	74.6	82.9	84.4
<i>SIXray_c</i>	FA [25]	77.1	84.5	85.8
<i>SIXray_c</i> \Rightarrow	CR-CNN [24]	73.8	83.7	84.5
<i>SIXray_c</i>	FA [25]	76.7	85.6	85.7

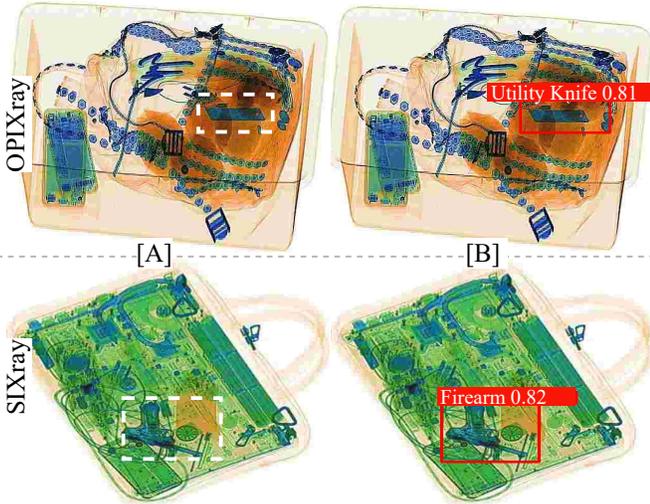


Fig. 4. Detection using FreeAnchor [25] at compression of 10: [A] trained on uncompressed, [B] retraining with compressed images. White box highlights the missing detection.

FreeAnchor [25] achieves mAP of 87.4 (*OPIXray_c*) when retrained at a compression level of 10%, nearly matching its performance of uncompressed settings, while affording a lossy compression rate much higher in terms of reduced image storage requirements. The impact of lossy compression at high levels is similarly mitigated through retraining with Cascade R-CNN [24], achieving a mAP of 84.9 when retrained at a level of 10%. For *SIXray*, retraining slightly improves upon the performance on both architectures at compression levels of 50% and 95%, reducing the mAP decrease to 0.4%, but does not alleviate the performance drop at 10%. The performance of both models suffer by 0.8% when retrained at a level of 10%, compared to the models trained on the uncompressed *SIXray*. This is likely due to the lower resolutions of images in *SIXray* [5] compared to *OPIXray* [4], such that the former are more adversely affected by additional compression in terms of overall image quality.

C. Effect of Data Augmentation

Table IV presents the results from evaluating CNN models with varying data augmentation techniques (Section II-B)

using *OPIXray* [4] and *SIXray* [5] respectively. To provide a baseline performance we train the models without data augmentation and use this to compare to each technique. All augmentations have an application probability of 0.5 except for *RandomCrop* that randomly selects a crop size between 75% and 100% of the image height and width dimensions. *JPEG* describes compressing the input image to a JPEG quality level of 10%. After examining the performance of *CutMix* we evaluate *ClassCutmix* that combines two selected classes when they are present. These are *straight knife* with *utility knife* for *OPIXray* [4] and *pliers* with *knife* for *SIXray* [5]. The values contained within brackets in Table IV denotes the change in mAP from the baseline value.

TABLE IV
EFFECT OF DATA AUGMENTATION ON MODEL PERFORMANCES (mAP).

Augmentation	<i>OPIXray</i>		<i>SIXray</i>	
	FA [25]	CR-CNN [24]	FA [25]	CR-CNN [24]
Baseline	87.7	85.1	85.2	84.3
RandomFlip	87.7 (+0.0)	85.9 (+0.9)	85.8 (+0.6)	84.6 (+0.3)
RandomCrop	90.9 (+3.2)	88.2 (+3.1)	84.4 (-0.8)	84.1 (-0.2)
Rotate	88.8 (+1.1)	86.0 (+0.9)	83.0 (-2.2)	83.5 (-0.8)
Equalise	86.5 (-1.2)	85.4 (+0.3)	84.7 (-0.5)	84.4 (+0.1)
Blur	87.8 (+0.1)	86.4 (+1.3)	85.3 (+0.1)	84.1 (-0.2)
JPEG	87.7 (+0.0)	86.1 (+1.0)	85.3 (+0.1)	84.1 (-0.2)
MixUp	88.5 (+0.8)	85.8 (+0.7)	84.2 (-1.0)	84.1 (-0.2)
BboxMixUp	88.0 (+0.3)	85.4 (+0.3)	84.0 (-1.2)	83.4 (-0.9)
CutMix	87.3 (-0.4)	85.6 (+0.5)	84.8 (-0.4)	82.9 (-1.4)
ClassCutMix	87.4 (-0.3)	86.3 (+1.2)	83.9 (-1.3)	82.5 (-1.8)

When applied to *OPIXray* [4] the data augmentations consistently improve Cascade R-CNN [24] performance but their impact is mixed for FreeAnchor [25]. *RandomCrop* produces the highest improvement for both models, increasing performance nearly identically by 3.65% and 3.64% for FreeAnchor [25] and Cascade R-CNN [24] respectively, with both exceeding their previous maximal performances (Table I) without a notable increase in model complexity. Demonstrating that complete occlusion of parts of the image (via cropping) is an effective method to generate further samples from the images within *OPIXray* [4]. MixUp (Section II-B) increases FreeAnchor [25] performance by 0.9% and Cascade R-CNN by 0.8%, demonstrating that it does indeed create beneficial new augmented images for *OPIXray* [4]. The bounding box variation of MixUp (*BboxMixUp*) generates a smaller improvement compared to full image *MixUp* of 0.35%, suggesting that with *OPIXray* [4] the models benefit from the increased background complexity, not just further prohibited items introduced into the image. Both variations of *CutMix* adversely affect the performance of FreeAnchor [25] model. In contrast, *CutMix* and *ClassCutmix* increased the performance of Cascade R-CNN [24] by 0.6% and 1.4% beyond its baseline. Applying the selected augmentation techniques to the *SIXray* [5] dataset results in the quality of models produced generally decreasing. Only *RandomFlip* improves the performance of both models, with FreeAnchor [25] achieving 85.8 mAP and Cascade R-CNN [24] achieving 84.6 mAP.

Unlike the *OPIXray* [4], the majority of images in *SIXray* [5] contain multiple prohibited items. This could limit the variation that can be introduced; items are already at multiple orientations and scales limiting geometric based augmenta-

tions, the range of image quality could reduce the impact of changing image properties, and techniques such as cropping or combining images could be more likely to include partial features. This also indicates that the *SIXray* [5] is a challenging dataset for prohibited item detection compared to *OPIXray* [4]. However, *OPIXray* [4] is a realistic dataset for X-ray security imagery applications, whereby an image primarily contains a single prohibited item, which more resembles a likely real-world scenario.

V. CONCLUSION

This work evaluates two CNN based object detection architectures on the task of threat item detection within X-ray imagery. FreeAnchor achieves superior performances (mAP: 87.7 and 85.8) using *OPIXray* [4] and *SIXray* [5] datasets. Both FreeAnchor [25] and Cascade R-CNN [24], outperform prior works [4], [5] for both datasets. Evaluation of the model complexity and runtime highlights the efficiency of the FreeAnchor [25] in terms of performance relative to number of model parameters, achieving a mAP:C ratio nearly $2\times$ that of Cascade R-CNN [24] across both datasets. We demonstrate retraining the networks on compressed imagery has been shown as a method to decrease the impact especially at higher levels of compression, which reduces the storage overhead for deployment within X-ray scanner applications. However, even a $1 \sim 2\%$ performance drop may not be acceptable in X-ray security screening involving safety-critical operation. A thorough evaluation of augmentation methods reveals a set of effective techniques, that assist in alleviating the performance limitations introduced by the low number of samples within existing prohibited item X-ray datasets. Future work will investigate the performance of varying backbone networks, such as ResNet₁₀₁ [29] in combination with different CNN object detection and segmentation architectures.

REFERENCES

- [1] "Gilardoni FEP-ME-640-AMX," www.gilardoni.it/en/security/x-ray-solutions/automatic-detection-of-explosives/fep-me-640-amx/, accessed: 2021-01-10.
- [2] M. Mendes, A. Schwaninger, N. Strelbel, and S. Michel, "Why laptops should be screened separately when conventional x-ray screening is used," in *Proc. Int. Carnahan Conf. on Security Technology*. IEEE, 2012, pp. 267–273.
- [3] S. Michel, S. M. Koller, J. C. de Ruyter, R. Moerland, M. Hogervorst, and A. Schwaninger, "Computer-based training increases efficiency in x-ray image interpretation by aviation security screeners," in *Proc. Int. Carnahan Conf. on Security Technology*, 2007, pp. 201–206.
- [4] Y. Wei, R. Tao, Z. Wu, Y. Ma, L. Zhang, and X. Liu, "Occluded prohibited items detection: An x-ray security inspection benchmark and de-occlusion attention module," in *Proc. Int. Conf. on Multimedia*, ser. MM '20, 2020, p. 138–146.
- [5] C. Miao, L. Xie, F. Wan, C. Su, H. Liu, J. Jiao, and Q. Ye, "SIXray: A large-scale security inspection x-ray benchmark for prohibited item discovery in overlapping images," in *Proc. on Computer Vision and Pattern Recognition*, 2019, pp. 2114–2123.
- [6] M. Kundegorski, S. Akcay, M. Devereux, A. Mouton, and T. Breckon, "On using feature descriptors as visual words for object detection within x-ray baggage security screening," in *Proc. Int. Conf. on Imaging for Crime Detection and Prevention*, November 2016, pp. 1–6.
- [7] D. Mery, V. Rizzo, I. Zuccar, and C. Pieringer, "Automated x-ray object recognition using an efficient search algorithm in multiple views," in *Proc. Int. Conf. on Computer Vision and Pattern Recognition Workshops*, 2013, pp. 368–374.
- [8] S. Akcay, M. E. Kundegorski, C. G. Willcocks, and T. P. Breckon, "Using deep convolutional neural network architectures for object classification and detection within x-ray baggage security imagery," *IEEE Transactions on Information Forensics and Security*, vol. 13, no. 9, pp. 2203–2215, 2018.
- [9] Y. Gaus, N. Bhowmik, S. Akcay, and T. Breckon, "Evaluating the transferability and adversarial discrimination of convolutional neural networks for threat object detection and classification within x-ray security imagery," in *Proc. Int. Conf. on Machine Learning Applications*, 2019, pp. 420–425.
- [10] N. Bhowmik, W. Q., Y. Gaus, M. Szarek, and T. Breckon, "The good, the bad and the ugly: Evaluating convolutional neural networks for prohibited item detection using real and synthetically composite x-ray imagery," in *Proc. British Machine Vision Conference Workshops*. BMVA, 2019, pp. 1–8.
- [11] S. Ren, K. He, R. Girshick, and J. Sun, "Faster R-CNN: Towards real-time object detection with region proposal networks," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 39, no. 6, pp. 1137–1149, 2017.
- [12] J. Dai, Y. Li, K. He, and J. Sun, "R-FCN: Object detection via region-based fully convolutional networks," in *Advances in Neural Information Processing Systems*, 2016, pp. 379–387.
- [13] J. Redmon and A. Farhadi, "Yolo9000: Better, faster, stronger," in *Proc. on Computer Vision and Pattern Recognition*, 2017, pp. 6517–6525.
- [14] T. Lin, P. Goyal, R. Girshick, K. He, and P. Dollár, "Focal loss for dense object detection," in *Proc. Int. Conf. on Computer Vision*, 2017, pp. 2999–3007.
- [15] S. Akçay, M. E. Kundegorski, M. Devereux, and T. P. Breckon, "Transfer learning using convolutional neural networks for object classification within x-ray baggage security imagery," in *Proc. Int. Conf. on Image Processing*, 2016, pp. 1057–1061.
- [16] M. Poyser, A. Atapour-Abarghouei, and T. Breckon, "On the impact of lossy image and video compression on the performance of deep convolutional neural network architectures," in *Proc. Int. Conf. Pattern Recognition*, 2020.
- [17] S. Dodge and L. Karam, "Understanding how image quality affects deep neural networks," in *Proc. Int. Conf. on Quality of Multimedia Experience*, 2016, pp. 1–6.
- [18] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "ImageNet: A large-scale hierarchical image database," in *Proc. Conf. Computer Vision and Pattern Recognition*, 2009, pp. 248–255.
- [19] M. Everingham, L. Van Gool, C. K. Williams, J. Winn, and A. Zisserman, "The pascal visual object classes (VOC) challenge," *Int. journal of computer vision*, vol. 88, no. 2, pp. 303–338, 2010.
- [20] P. Y. Simard, Y. A. LeCun, J. S. Denker, and B. Victorri, *Transformation Invariance in Pattern Recognition — Tangent Distance and Tangent Propagation*. Springer Berlin Heidelberg, 1998, pp. 239–274.
- [21] C. Shorten and T. Khoshgoftaar, "A survey on image data augmentation for deep learning," *Journal of Big Data*, vol. 6, 07 2019.
- [22] S. Yun, D. Han, S. Chun, S. J. Oh, Y. Yoo, and J. Choe, "Cutmix: Regularization strategy to train strong classifiers with localizable features," in *Proc. on Int. Conf. on Computer Vision*, 2019, pp. 6022–6031.
- [23] H. Zhang, M. Cisse, Y. N. Dauphin, and D. Lopez-Paz, "mixup: Beyond empirical risk minimization," in *Proc. on Int. Conf. on Learning Representations*, 2018.
- [24] Z. Cai and N. Vasconcelos, "Cascade R-CNN: High quality object detection and instance segmentation," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pp. 1–1, 2019.
- [25] X. Zhang, F. Wan, C. Liu, R. Ji, and Q. Ye, "Freeanchor: Learning to match anchors for visual object detection," in *Advances in Neural Information Processing Systems*, vol. 32, 2019, pp. 147–155.
- [26] G. K. Wallace, "The jpeg still picture compression standard," *Commun. ACM*, vol. 34, no. 4, p. 30–44, Apr. 1991.
- [27] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick, "Microsoft COCO: Common objects in context," in *Computer Vision – ECCV 2014*, 2014, pp. 740–755.
- [28] K. Chen, J. Wang, J. Pang, Y. Cao, Y. Xiong, X. Li, S. Sun, W. Feng, Z. Liu, J. Xu, Z. Zhang, D. Cheng, C. Zhu, T. Cheng, Q. Zhao, B. Li, X. Lu, R. Zhu, Y. Wu, J. Dai, J. Wang, J. Shi, W. Ouyang, C. C. Loy, and D. Lin, "MMDetection: Open mmlab detection toolbox and benchmark," *arXiv preprint arXiv:1906.07155*, 2019.
- [29] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. on Computer Vision and Pattern Recognition*, 2016, pp. 770–778.