

Rotational Rectification Network: Enabling Pedestrian Detection for Mobile Vision

Xinshuo Weng
Carnegie Mellon University
xinshuow@andrew.cmu.edu

Shangxuan Wu
Carnegie Mellon University
shangxuw@andrew.cmu.edu

Fares Beainy
Volvo Construction Equipment
fares.beainy@volvo.com

Kris M. Kitani
Carnegie Mellon University
kkitani@cs.cmu.edu

Abstract

Across a majority of pedestrian detection datasets, it is typically assumed that pedestrians will be standing upright with respect to the image coordinate system. This assumption however, is not always valid for many vision-equipped mobile platforms such as mobile phones, UAVs or construction vehicles on rugged terrain. In these situations, the motion of the camera can cause images of pedestrians to be captured at extreme angles. This can lead to very poor pedestrian detection performance when using standard pedestrian detectors. To address this issue, we propose a Rotational Rectification Network (R2N) that can be inserted into any CNN-based pedestrian (or object) detector to adapt it to significant changes in camera rotation. The rotational rectification network uses a 2D rotation estimation module that passes rotational information to a spatial transformer network [12] to undistort image features. To enable robust rotation estimation, we propose a Global Polar Pooling (GP-Pooling) operator to capture rotational shifts in convolutional features. Through our experiments, we show how our rotational rectification network can be used to improve the performance of the state-of-the-art pedestrian detector under heavy image rotation by up to 45%.

1. Introduction

Pedestrian detection is an active research area in computer vision and has rapidly progressed through the recent decade. There are many benchmark pedestrian detection datasets for learning and evaluation [6, 10, 7, 4]. One common setting in these datasets is that the camera’s y-axis is roughly aligned to the direction of gravity. This means that pedestrians are captured in the vertical direction because pedestrians usually stand upright on the ground. This




Figure 1. **Upper:** Schematic diagram of how the proposed rotation rectification network (blue) is inserted into a CNN-based pedestrian detector (yellow). **Lower:** Sample detection results produced by a state-of-the-art pedestrian detector (**left**) and with the R2N inserted (**right**). Inserting the R2N increases the robustness of the network to image rotation and decreases the miss rate of detection.

“upright assumption” in benchmark datasets distinguishes pedestrians from many objects in the scene. Much work has been devoted to designing features [4, 9] or model architectures [29, 21, 25] to learn the appearance of upright pedestrians.

However, this upright assumption may not always be true in real-world situations where the camera orientation can be highly dynamic. For example, when recording a video with a mobile phone camera, the angle of the camera can vary wildly as one walks or runs. For cameras installed on construction vehicles, the upright assumption is easily invalidated when recording video over rough terrain. In both of these examples, the projection of pedestrians in the image

can be at extreme angles of rotation. Detecting pedestrians in such situations is difficult with current state-of-the-art models.

One straightforward way to achieve rotational robustness for pedestrian detection is to simply increase the size of the training data to include more instances of pedestrians imaged at an angle. When new data is not available, one can also augment the data [26] by creating new training examples from rotated images. While such a brute-force approach can certainly lead to improvements, simply creating more data does not necessarily address the fundamental problem of understanding and modeling image rotations.

An alternative solution is to attempt to infer the rotational distortion of the image and to remove the effect of that distortion prior to detection within a unified rotation-invariant detection network. However, estimating rotational changes in an image is difficult with the current paradigm of convolutional feature extraction because they are based on a rectangular spatial decomposition of the image. In other words, rotational changes in image content can produce very different feature responses in the upper convolutional layers of a convolutional neural network (CNN).

In order to facilitate a smoother change in convolutional feature responses due to image rotation, we propose the use of a novel Global Polar Pooling (GP-Pooling) operator which converts rectangular convolutional feature responses to a polar grid system. Using this polar coordinate system, rotations of the input images only result in translational shifts of the polar features maps, making it easier for higher level convolution layers to model image rotation. In this way, our proposed GP-Pooling operator gives a deep neural network the ability to encode image rotations more effectively.

In order to obtain rotational invariance during detection we propose a rotational rectification network (R2N) that can be flexibly inserted into an intermediate layer of a general detection network. The R2N uses a CNN with a GP-Pooling layer to estimate the rotation angle present in the images. Then the estimated rotation θ is passed to a spatial transformer network to undistort the image features. An overview of the overall network architecture is illustrated in Figure 2. We show that after removing the effect of rotation inside a network, the general detector can be easily adapted to work on pedestrians imaged at arbitrary rotation angles.

The main contributions of our work are summarized as follows: (1) we propose a Global Polar Pooling (GP-Pooling) operator, which can be used to encode the radial distribution of features within a general CNN architecture and (2) we propose a rotational rectification network (R2N) that can be inserted into a wide range of CNN-based detectors to achieve rotational invariance.

2. Related Work

Rotational Robustness in CNNs. The existing methods to incorporate CNN with rotational robustness can be split into two categories. The first category methods add robustness by manipulating the images or feature maps. Prior works [1, 14, 11, 8, 13] achieve rotational robustness by augmenting the input images on the fly and fusing the response in the upper layer of the network. Dieleman *et al.* [5] copy the feature responses from intermediate layers in four 90° angles and compress them by their proposed operators, which achieves rotational robustness in a compact way. However, by augmenting the data, these methods are only robust to a discrete set of rotations instead of 360° continuous rotations. This does not directly address the fundamental problem of incorporating rotation rectification in CNNs.

[12] introduce a general warp framework called Spatial Transformer Networks to enable affine transformations with differentiable sampling inside the network. It achieves transformation invariance within CNN architectures very efficiently without data augmentation. An important point which is often overlooked is that the design of the *localization network* whose purpose is to estimate transformation parameters is not explored in the original work. They use two baseline CNNs as the localization networks in the spatial transformer: (1) two fully-connected layers; (2) a CNN with two convolutional and two max pooling layers. The models are evaluated on the distorted MNIST dataset. However, the MNIST dataset is small and has low resolution images which is not a strict criteria to judge transformation invariance of a network. In other words, the design of a network with natural transformation invariance is still an open problem. Note that our work is complementary to the spatial transformer because our proposed rotation estimation module with GP-Pooling operators can be viewed as an expert localization network with natural rotation invariance.

On the other hand, the second category methods achieve rotational robustness by modifying filters within CNN architectures instead of manipulating the data. Cohen *et al.* [2, 3] apply kernel-based pooling to sample responses in symmetry space such that only the least important features are lost at each layer. Prior works [19, 30, 20, 22] replicate and transform the learned canonical filters in a finite set of orientations and then fuse the output responses at each layer to achieve rotational robustness. However, these methods are also robust to only a discrete set of rotations. Instead, our proposed GP-Pooling operator is able to add rotational robustness to general CNN architectures in 360° continuous rotations.

H-Nets [27] replace regular CNN filter with complex circular harmonics and is also able to capture continuous rotational changes. But H-Nets assume the learned filters are in the harmonic wavelets space. Instead, GP-Pooling

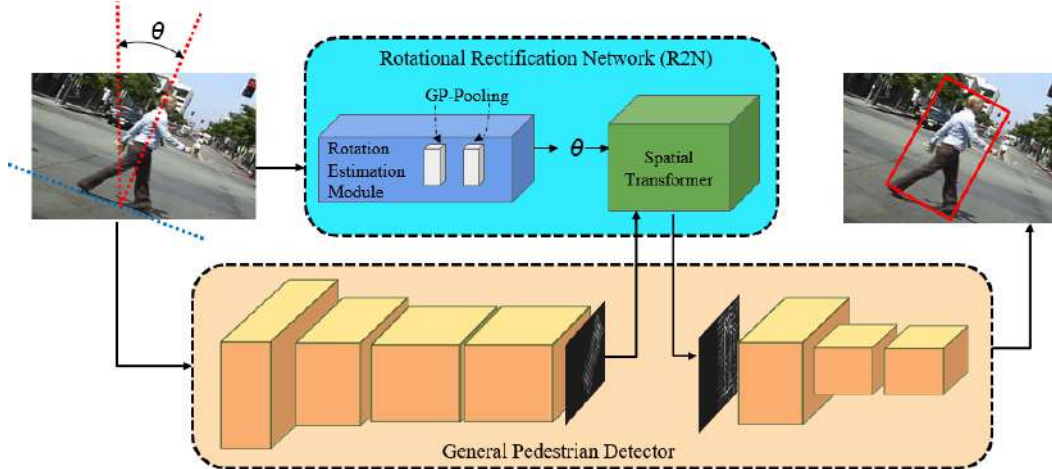


Figure 2. Architectural Overview. Rotation rectification network (R2N) (cyan) is inserted into intermediate layer of pre-existing CNN-based pedestrian detector (yellow). R2N uses rotation estimation network (see Figure 4) with GP-Pooling (gray) operators to estimate rotation angle (blue). Estimated rotation angle θ is passed to Spatial Transformer (green). R2N warps image features to remove global rotation. Last layer (yellow) yields tight rotated bounding boxes.

does not impose any assumption on the image filters. Moreover, H-Nets is designed to learn local rotational robust filters while GP-Pooling operator focuses on global rotational changes. More importantly, most existing methods only test rotational robustness on synthetic tasks such as digit recognition from MNIST. Our proposed GP-Pooling operator is able to succeed on real-world task, namely, pedestrian detection on the Caltech Pedestrians dataset.

Detection More recent detection methods are based on region proposals. They perform detection by classifying region proposal of images and regressing the bounding box simultaneously. For example, Ren *et al.* [23] introduce a Region Proposal Network (RPN) to enable nearly cost-free region proposals and propose an unified detection framework. Liu *et al.* [17] introduce default boxes which tiles input images and then regress the offset for each box in the work of Single Shot MultiBox Detector (SSD). In the context of pedestrian detection, Zhang *et al.* [29] analyze the performance of Faster-RCNN on pedestrian detection and propose a simple yet powerful baseline for pedestrian detection based on RPN. In many of these methods, the region proposals are represented by axis-aligned rectangles, which are not suitable for detecting pedestrians imaged at an angle. To address this issue, Ma *et al.* [18] propose a novel framework to detect text with arbitrary orientation in natural scene images. In their work, they present the Rotation Region Proposal Networks (RRPN) to generate rectangular proposals at different rotations instead of axis-aligned proposals only. This approach is limited because the RRPN can only deal with a discrete set of rotations and it is only applicable to proposal-based detection network.




Figure 3. **Left:** GP-Pooling Operator. Feature map in polar coordinates. **Right:** Feature Responses of GP-Pooling Operator. Input rotations result in translational shifts of feature responses.

3. Global Polar Pooling (GP-Pooling)

In a CNN, pooling increases the receptive field and filters out the noisy feature responses from previous layers. Moreover, existing pooling operators, especially spatial max pooling, are used frequently because of their robustness to translation. In other words, translational changes in image content can produce feature responses with translational shifts in existing spatial pooling operators. However, this is not the case for rotational changes. Current paradigms of convolutional feature extraction are strictly based on a rectangular spatial decomposition of the image features. As a result, any rotational changes in image content can produce a very different feature response.

To achieve a smoother change in general CNN feature responses due to rotational changes, we propose the Global Polar Pooling (GP-Pooling) operator. This operator extends existing pooling operators from rectangular to a radial decomposition. It makes the global rotational changes from the input image content easily recognizable in CNNs. Specifically, the GP-Pooling operator represents convolutional feature responses on a polar grid system such that any

in plane rotation from an input image results in translational shift of the polar feature map. Then the translational shifts can be used by the upper layers of the network to enable rotation invariance.

The core idea of how GP-Pooling works is illustrated in Figure 3. Inside the GP-Pooling operator, the feature maps are represented in polar coordinate system. The origin is defined at the center of the feature map. We note here that while our pooling layer is designed primarily for in-plane rotation about the image center, we find empirically that it can handle moderate levels of off center rotation. To be concrete, a pixel $P(x, y)$ of the feature map with width w and height h can be represented in polar coordinate $P_p(x_p, y_p)$ by:

$$P_n(x_n, y_n) = \left(x - \frac{w}{2}, -y + \frac{h}{2}\right) \quad (1)$$

$$P_p(x_p, y_p) = (\sqrt{x_n^2 + y_n^2}, \text{atan2}(y_n, x_n)), \quad (2)$$

where $P_n(x_n, y_n)$ is the normalized coordinate based on the center of the feature map.

The key difference of GP-Pooling from existing pooling operators is that we define parameters of kernel size, stride and padding along radial and angular axis in polar coordinates. These parameters determine how the input feature map is tiled into a grid. Inside each grid cell, a max operation is then executed for pixels that fall into that cell. A switch variable records the location of maximum activation. Then, the gradient only flows back to this location during backpropagation. This is illustrated in Figure 3, where the input feature map is tiled with kernel size of $\frac{\pi}{4}$, stride of $\frac{\pi}{4}$ and padding of 0 along angular axis. In this particular illustrative example, this results in 8 angular sectors, each of size $\frac{\pi}{4}$ radians. Each angular sector is further divided into 7 cells along radial axis. In practice, it is necessary to set the stride and kernel size along the angular axis to be $\frac{\pi}{180}$ to capture one degree of image rotation.

To demonstrate the functionality of converting rotational changes to translational shifts, we visualize the output features of the proposed GP-Pooling operator in Figure 3. We take two MNIST images (28×28) as the input of the GP-Pooling operator. The kernel size, stride and padding along angular and radial axis is $\frac{\pi}{36}$, $\frac{\pi}{36}$, 0, 1, 1 and 0 respectively. This results in an output features of size 20×72 . We then re-scale them to a size of 28×100 to obtain the same height as the input image for better visualization. The results show that the feature responses approximately shift leftwards or rightwards when we rotate the input image.

Rotation Estimation Module. In order to estimate the rotation parameter from images or image features efficiently, the property of converting rotational changes to translational shifts is very useful. We insert multiple GP-Pooling layers into the rotation estimation module in a multi-scale manner and concatenate their output feature responses with features from general spatial max pooling layer. The rotation estimation module ultimately outputs the estimation of

rotation θ ranging from $-\pi$ to π present in the input image by solving a regression problem. A typical architecture of our rotation estimation module embedded with GP-Pooling operator is shown in Figure 4.

4. Rotational Rectification Network (R2N)

We describe the rotational rectification network (R2N) and how we fit it into a general pedestrian detector to achieve rotation-robust detection. The R2N takes two inputs: (1) the input image and (2) an intermediate feature map from an intermediate layer of the detection network, then it outputs a warped image feature where the global rotation has been removed. In essence, we transform a complex task of arbitrary-oriented pedestrian detection to a much easier task of upright pedestrian detection. The overall architecture is shown in Figure 2. The R2N is composed of a rotation estimation module and a spatial transformer. We use the estimation of rotation θ from the rotation estimation module to construct a 2×3 transformation (rotation) matrix M .

$$M = \begin{bmatrix} \cos\theta & -\sin\theta & 0 \\ \sin\theta & \cos\theta & 0 \end{bmatrix} \quad (3)$$

The transformation matrix M is then passed to the spatial transformer to warp the input feature map. This warping removes the effect of global rotation distortion present in the image features prior to detection. Note that in our paradigm, the spatial transformer is used only for image warping. The warp parameters are provided by our specific localization net¹ – the rotation estimation module – which is designed for increasing robustness to image rotation.

We emphasize here again, that the R2N is a separate module, independent of the pedestrian detector. It can work as a plugin and be inserted into an intermediate layer in many CNN-based pedestrian detectors to achieve rotation-invariant detection. In practice, we usually insert the R2N module into the feature extraction part of the network (e.g., after the pool3 layer of the VGG part). The R2N can enable the final layer of pedestrian detector to yield the tight rotated bounding boxes based on the estimated rotation θ .

5. Datasets

In order to evaluate the performance of our GP-pooling layer and our rotation invariant R2N network, we need a dataset with images of people undergoing heavy rotation. In order to perform a detailed quantitative analysis, we utilize a rotated MNIST and a rotated Caltech datasets, where the digit and pedestrian images are synthetically rotated at various angles. In order to verify the performance on real rotated data in the wild, we collected a YouTube Wearable

¹Please refer to [12] for more details about the spatial transformer and definition of localization net.




Figure 4. Architecture of a rotation estimation module embedded with the GP-Pooling (blue) operator. This network composes of convolution (red), max pooling (yellow), GP-Pooling (blue), batch normalization (gray), concatenation (green), flatten (magenta), fully connected layer (cyan). The rotation estimation module takes images or image features as the input. The final regression layer produces the estimation of rotation θ present in the input.

Video dataset, where images are captured by people with wearable cameras during dynamic activities (e.g., running, riding a bike) such that pedestrians are imaged at various angles.

5.1. Rotated MNIST

The rotated MNIST dataset is created by rotating images from the MNIST dataset [16]. The rotation angle is uniformly selected from -90° to 90° (the upper half of a circle). The rotated MNIST dataset contains 10000 training images, 2000 validation images and 50000 testing images, each with a size of 28×28 . We emphasize here that this dataset is not used to evaluate object classification but rather to evaluate the performance of image rotation estimation.



Figure 5. Sample images from the rotated MNIST dataset.

5.2. Rotated Caltech

We use a rotated version of the Caltech pedestrian dataset to evaluate the ability of a pedestrian detector to detect people imaged at varying angles of rotation. The original Caltech dataset contains 6 video sequences for training and 5 for testing. Since consecutive images are very similar, images are sampled every 3 frames in training set and every 30 frames in testing set. This results in 42786 training images and 4024 testing images in our rotated Caltech dataset. We rotate all images by uniformly selecting the rotation angle from -90° to 90° . The rotation angles are saved as ground truth for training and testing the rotation estimation module.



Figure 6. Sample images from the rotated Caltech dataset.

5.3. YouTube Wearable

In order to evaluate the performance of our rotation invariant R2N in the wild, we create the YouTube Wearable dataset where images contain pedestrians with various poses without manual rotation. We obtain 100 short YouTube videos recorded by people with wearable cameras.

To have a high possibility to contain pedestrians within images, we cut 4000 frames from videos where images are taken outdoor in the city. Around 500 images contain pedestrians. We label the bounding boxes manually and remove boxes under 400 pixels.



Figure 7. Sample images from the YouTube Wearable dataset.

6. Evaluating the Rotation Estimation Module

The success of the R2N module to enable rotation invariant detection relies heavily on the precision of the estimated rotation θ . As such, it is critical to estimate the image rotation precisely. In this first experiment, we evaluate the accuracy of the rotation estimation module using the proposed GP-Pooling operator on rotated MNIST and rotated Caltech dataset.

6.1. Rotation Estimation on Rotated MNIST

In order to see how the rotation estimation module works independently despite of the R2N, we evaluate three baseline rotation estimation modules. Descriptions of the baselines are in Table 1 (first three rows). We train all baselines using rotation angle as the ground truth and evaluate them based on root of sum of squared error.

Training Details. we use Euclidean loss during training. The training takes 160 epochs with batch size of 128 for all models. Adam optimizer is used with learning rate of 0.001, two momentums of 0.9 and 0.999.

Results. In Table 3, we find that the 2-layer fully connected network cannot work very well because the network has too few neurons and simple structure. Importantly, the GP-Pooling+CNN outperforms the CNN by achieving 22% lower error (from 11.26° to 8.78°), showing that the adding GP-Pooling operator to the network is able to increase robustness to image rotation.

Visualization. To understand what has been learned in the rotation estimation module (GP-Pooling+CNN), we extract the feature responses before the fc3 layer as a vector representation for each image patch. For a set of 7 random digit

Name	With STN?[12]	Description
2FC	No	a 2-layer fully connected network (20 and 1 neurons per layer respectively).
CNN	No	a standard CNN described in Table 2.
CNN+GP-Pooling	No	Two GP-Pooling operators are inserted into the CNN in 2nd row after the conv1 and conv2 layers before concatenating with the pool2 features.
STN-2FC [12]	Yes	a spatial transformer with 2FC as the rotation estimation module.
STN-CNN [12]	Yes	a spatial transformer with CNN as the rotation estimation module.
STN-CNN+GP-Pooling	Yes	a spatial transformer with CNN+GP-Pooling as the rotation estimation module.

Table 1. Baselines used in two rotated MNIST experiments.

Layer	conv1	ReLU	pool1	conv2	ReLU	pool2	fc3	ReLU	dropout	fc4	sigmoid
Units	16	16	16	32	32	32	20	20	20	1	1
Feature	28×28	28×28	14×14	14×14	14×14	7×7	1	1	1	1	1

Table 2. Topology of the CNN used in the rotated MNIST dataset.

images, we find the 19 nearest neighbors (Figure 8). We observe that the nearest neighbors are *not* necessarily the same digit. More importantly, the nearest neighbors have the same rotation angle as the query image which seems to indicate that the learned feature representation is encoding digit angle instead of digit label.

6.2. Rotation Estimation with Spatial Transformer Network (STN) on Rotated MNIST

As our proposed rotation estimation module with the GP-Pooling operator is similar but complimentary to the spatial transformer (STN), we evaluate the STN with/without our GP-Pooling to see if the GP-Pooling is helpful in STN. we use the same rotation estimation modules evaluated in section 6.1. Descriptions of the baselines are in Table 1 (**last three rows**). As we estimate the rotation parameters, the transformation matrix of the STN is restricted to a rotation matrix in the form of Equation 3. We train three baselines STNs using digit labels as the ground truth because STN is trained with digit classification task in MNIST dataset. However during testing, in order to see how the GP-Pooling operator affects the performance of rotation estimation, we do not care about the predicted digit but evaluate the estimation of rotation produced by the rotation estimation module, an intermediate output of the overall network.

Training Details. As we are solving a classification task during training, we use Cross Entropy Loss instead of Euclidean Loss. All other training details are the same as in section 6.1.

Results. In table 4, we observe that, while STNs are trained with classification task without using rotation angle, the rotation estimation module is learning to estimate the image rotation. A stronger rotation estimation module (STN-CNN) can have lower estimation error than a simpler model (STN-2FC). Additionally, we found that adding the GP-Pooling operators to the rotation estimation module helps improve the accuracy of rotation estimation even with the

Methods	Error (degree)
2FC	44.91°
CNN	11.26°
GP-Pooling+CNN	8.78°

Table 3. Rotation estimation error on rotated MNIST dataset. 2FC, CNN and GP-Pooling+CNN are the rotation estimation modules defined in Table 1.

Methods	Error (degree)
STN-2FC [12]	23.37°
STN-CNN [12]	18.00°
STN-GP-Pooling+CNN	16.38°

Table 4. Rotation estimation error on Rotated MNIST dataset with Spatial Transformer Network.




Figure 8. Feature Space Visualization. Nearest neighbors of probe images (left column). Feature representation encodes digit angle, not digit label.

spatial transformer.

6.3. Rotation Estimation on Rotated Caltech

As the rotated MNIST dataset is very simple and has only low resolution images, it can not be strict dataset to judge the rotation invariance of a network. As such, we also evaluate the rotation estimation module on rotated Caltech dataset, where images have much higher resolution (480 × 640) and more complex contents compared to the MNIST. However, the task of pedestrian detection on Caltech dataset is beyond the scope of the spatial transformer, so the comparison experiments with STN are not applicable here. On the other hand, the task of horizontal line detection is very

Methods	Error (degree)
Zhai et al [28]	38.79°
Lezama et al [15]	29.26°
VGG-S [24]	18.33°
GP-Pooling+VGG-S	15.82°

Table 5. Results on rotated Caltech dataset. The topology of the GP-Pooling+VGG-S is described in Figure 4 and Section 5.2.

similar to the rotation estimation (estimation of rotation is the slope of horizontal line), so we compare our rotation estimation module with two state-of-the-art horizontal line detection algorithms.

Baselines.

1. **Zhai et al. [28]**: a state-of-the-art CNN-based horizontal line detector.
2. **Lezama et al. [15]**: a traditional edge-based horizontal line detector.
3. **VGG-S [24]**: the small version of VGG network.
4. **GP-Pooling+VGG-S**: two GP-Pooling operators are inserted to VGG-S after the conv2 and conv5 layers before concatenation. The topology of the network is shown in Figure 4. All convolution layers have kernel size of 3, stride of 1 and padding of 1. All max pooling layers have kernel size of 2 and stride of 2.

Training Details. For baseline 1 and 2, we follow the training procedure from the original work. For baseline 3 and 4, We fine-tune both networks from the VGG-S model pre-trained on ImageNet (up to the pool5 layer) for 16 epochs with batch size of 4, learning rate of 0.000001 and weight decay of 0.00005.

Results. In Table 5, we observe that [29] and [15] is not working as well as VGG-S and GP-Pooling+VGG-S. This is because both two horizontal line detection algorithms heavily rely on the geometric priors which are not always true on the rotated Caltech dataset, especially when the image rotation is heavy. Importantly, when comparing GP-Pooling+VGG-S with VGG-S, we can achieve lower error in degree by 13.7% (from 18.3° to 15.8°) by simply adding two GP-Pooling operators to the rotation estimation module. This demonstrates again that the proposed GP-Pooling operator can increase robustness to image rotation and improve rotation estimation.

7. Evaluating the Rotational Rectification Network (R2N)

In order to see our R2N can enable rotation invariance in CNN-based pedestrian detectors, we now evaluate the performance of two end-to-end pedestrian detectors with R2N on original/rotated Caltech and YouTube Wearable dataset.

Baselines. Faster-RCNN² [23] and RPN-BF³ [29] are very

²The Faster-RCNN we used is pre-trained on VOC2007 dataset. We only evaluate the class of “person” from the total 20 classes in VOC2007.

³For RPN-BF, we have not used the random forest part compared to the

strong pedestrian detectors, we use four variants of them as the baselines for comparative analysis:

1. **Base**⁴: Faster-RCNN is trained on VOC2007 dataset and RPN-BF is trained on original Caltech dataset.
2. **Base+Aug**: Faster-RCNN and RPN-BF fine-tuned on the mixture of original and rotated Caltech dataset for better domain adaptation.
3. **Base+R2N**: We train the Faster-RCNN on VOC2007 dataset and RPN-BF on the original Caltech dataset, and then insert the proposed R2N module after the pool3 layer of both detectors without fine-tuning on rotated Caltech dataset.
4. **Base+R2N+GT**: Instead of estimating the rotation from the rotation estimation module, we pass the ground truth of rotation to the STN such that all pedestrians within images lead to upright pose. This should achieve upper-bound performance of the Base+R2N.

The evaluation metric for all following experiments is the log-average miss rate on false positive per image (FPPI) [6]. As is common practice, an intersection over union (IoU) of 0.5 is used to determine true positives.

Results on Rotated Caltech. We now evaluate the ability of our proposed R2N to transform a pre-trained pedestrian detector into a rotation-invariant pedestrian detector. We begin with experiments on rotated Caltech dataset. Results are shown in Figure 9 (**left**). The average miss rate (FPPI) of the base detectors (Faster-RCNN and RPN-BF) have a maximum of 89.8% and 83.3%, respectively, because both detectors are only trained on datasets without rotated pedestrians. By fine-tuning the detectors on a mixture of original and rotated Caltech dataset, the performance of Faster-RCNN and RPN-BF is increased 1% and 20% respectively. This is expected as better data augmentation is known to improve performance.

If we add the R2N to each detector, the performance increases by a large percentage: 11.0% for Faster-RCNN and 44.6% for RPN-BF (the average miss rate decreased from 83.3% to 57.6%). This result demonstrates that adding rotation invariance via our R2N module to a CNNs is much more effective at improving detection performance than data augmentation. We emphasize here, that we do not fine-tune the two baseline models after adding the R2N module. This shows that the R2N module works like a plugin without additional tuning.

Additionally, when comparing baseline base+R2N with base+R2N+GT, the performance is very close. This demonstrates that, in the case of no fine-tuning, the performance of base+R2N almost achieve the upper bound on rotated Caltech dataset. Qualitative results are shown in Figure 10.

original work.

⁴Base represents one of the general pedestrian detectors: Faster-RCNN or RPN-BF.




Figure 9. Quantitative results on rotated Caltech (**left**), original Caltech (**middle**) and YouTube Wearable dataset (**right**). Intersection of Union (IoU) of 0.5 is used to determine true positives.




Figure 10. Detection Results. Top 2 rows: rotated Caltech dataset. Bottom 2 rows: YouTube Wearable dataset.

Results on Original Caltech. We want to see how the added R2N module effects the performance of original pedestrian detector under upright images. Therefore, the second experiment focus on original dataset instead of a rotated one. When evaluated on the original Caltech dataset (Figure 9 **middle**), we found that the performance drops 1% and 3% for RPN-BF and Faster-RCNN respectively after inserting the R2N to the original detector. This drop is reasonable because we do not jointly fine-tune the networks after inserting the R2N module and the estimation of rotation from the rotation estimation module is not perfectly precise (i.e. it might add some rotations to the upright pedestrians.). Note that the performance of Faster-RCNN series is much worse than RPN-BF series due to the fact that Faster-RCNN is pre-trained on VOC2007 dataset not Caltech.

Results on YouTube Wearable. In order to see how our R2N performs on real-world rotated image data in the wild, we compare the Faster-RCNN and RPN-BF with/without our R2N on the YouTube Wearable dataset. Quantitative

and qualitative results are shown in Figure 9 (**right**) and 10 respectively. Compared to the results on Caltech dataset, Faster-RCNN+R2N performs much better, which might result from similar appearance and scale of pedestrians between YouTube Wearable and VOC2007 dataset. More importantly, adding the R2N increases the performance by 18.2% for Faster-RCNN and 11.7% for RPN-BF although we do not fine-tune the detectors on this completely new dataset. This demonstrates again that the proposed R2N can add rotation invariance to a detection network immediately without joint fine-tuning such that the detector is able to detect pedestrians with various poses.

8. Conclusion

We introduce the GP-Pooling operator which converts rotational changes to translational shifts and thus enables CNNs to encode rotational information. We then propose a rotational rectification network (R2N) and apply it to a real application of oriented pedestrian detection. We show

that the use of R2N can immediately help achieve rotation invariance without any fine-tuning given a detector trained on datasets with only upright pedestrians. Our approach enhance the performance of a state-of-the-art detector under heavy image rotation by 44.6% on rotated Caltech dataset.

References

- [1] G. Cheng, P. Zhou, and J. Han. Rfd-cnn: Rotation-invariant and fisher discriminative convolutional neural networks for object detection. *CVPR*, 2016.
- [2] T. Cohen and M. Welling. Group equivariant convolutional networks. *ICML*, 2016.
- [3] T. Cohen and M. Welling. Steerable cnns. *ArXiv e-prints*, 2017.
- [4] N. Dalal and B. Triggs. Histograms of oriented gradients for human detection. *CVPR*, 2005.
- [5] S. Dieleman, J. Fauw, and K. Kavukcuoglu. Exploiting cyclic symmetry in convolutional neural networks. *ArXiv e-prints*, 2016.
- [6] P. Dollár, C. Wojek, B. Schiele, and P. Perona. Pedestrian detection: An evaluation of the state of the art. *TPAMI*, 2012.
- [7] A. Ess, B. Leibe, K. Schindler, , and L. Gool. A mobile vision system for robust multi-person tracking. *CVPR*, 2008.
- [8] B. Fasel and D. Perez. Rotation-invariant neoperceptron. *ICPR*, 2006.
- [9] P. Felzenszwalb, R. Girshick, D. Mcallester, and D. Ramanan. Object detection with discriminatively trained part based models. *TPAMI*, 2009.
- [10] A. Geiger, P. Lenz, and R. Urtasun. Are we ready for autonomous driving? the kitti vision benchmark suite. *CVPR*, 2012.
- [11] J. Henriques and A. Vedaldi. Warped convolutions: Efficient invariance to spatial transformations. *ArXiv e-prints*, 2016.
- [12] M. Jaderberg, K. Simonyan, A. Zisserman, and K. Kavukcuoglu. Spatial transformer networks. *NIPS*, 2015.
- [13] D. Laptev and J. Buhmann. Transformation-invariant convolutional jungles. *CVPR*, 2015.
- [14] D. Laptev, N. Savinov, J. Buhmann, and M. Pollefeys. Tipooling: transformation-invariant pooling for feature learning in convolutional neural networks. *CVPR*, 2016.
- [15] J. Lezama, R. Gioi, G. Randall, and J. Morel. Finding vanishing points via point alignments in image primal and dual domains. *CVPR*, 2014.
- [16] C. Liu, K. Nakashima, H. Sako, and H. Fujisawa. Handwritten digit recognition: Benchmarking of state-of-the-art techniques. *CVPR*, 2003.
- [17] W. Liu, D. Anguelov, D. Erhan, C. Szegedy, and S. Reed. Ssd: Single shot multibox detector. *ArXiv e-prints*, 2015.
- [18] J. Ma, W. Shao, H. Ye, L. Wang, H. Wang, Y. Zheng, and X. Xue. Arbitrary-oriented scene text detection via rotation proposals. *ArXiv e-prints*, 2017.
- [19] D. Marcos, M. Volpi, N. Komodakis, and D. Tuia. Rotation equivariant vector field networks. *ArXiv e-prints*, 2016.
- [20] D. Marcos, M. Volpi, and D. Tuia. Learning rotation invariant convolutional filters for texture classification. *ArXiv e-prints*, 2016.
- [21] W. Ouyang and X. Wang. Joint deep learning for pedestrian detection. *ICCV*, 2013.
- [22] E. Oyallon and S. Mallat. Deep roto-translation scattering for object classification. *CVPR*, 2015.
- [23] S. Ren, K. He, R. Girshick, and J. Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. *NIPS*, 2015.
- [24] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. *ICLR*, 2015.
- [25] Y. Tian, P. Luo, X. Wang, and X. Tang. Deep Learning Strong Parts for Pedestrian Detection. *ICCV*, 2015.
- [26] D. van Dyk and X. Meng. The art of data augmentation. *JCGS*, 2001.
- [27] D. Worrall, S. Garbin, D. Turmukhambetov, and G. Brostow. Harmonic networks: Deep translation and rotation equivariance. *ArXiv e-prints*, 2016.
- [28] M. Zhai, S. Workman, and N. Jacobs. Detecting vanishing points using global image context in a non-manhattan world. *CVPR*, 2016.
- [29] L. Zhang, L. Lin, X. Liang, and K. He. Is faster r-cnn doing well for pedestrian detection? *ECCV*, 2016.
- [30] Y. Zhou, Q. Ye, Q. Qiu, and J. Jiao. Oriented response networks. *ArXiv e-prints*, 2017.