

Möbius Transform for Mitigating Perspective Distortions in Representation Learning

Prakash Chandra Chhipa¹, Meenakshi Subhash Chippa¹, Kanjar De², Rajkumar Saini¹, Marcus Liwicki¹, and Mubarak Shah³

¹ Luleå Tekniska Universitet, Sweden

{prakash.chandra.chhipa, rajkumar.saini, marcus.liwicki}@ltu.se
meechi-2@student.ltu.se

² Fraunhofer Heinrich-Hertz-Institut, Berlin, Germany

kanjar.de@hhi.fraunhofer.de

³ University of Central Florida, USA

shah@crcv.ucf.edu

Abstract. Perspective distortion (PD) causes unprecedented changes in shape, size, orientation, angles, and other spatial relationships of visual concepts in images. Precisely estimating camera intrinsic and extrinsic parameters is a challenging task that prevents synthesizing perspective distortion. Non-availability of dedicated training data poses a critical barrier to developing robust computer vision methods. Additionally, distortion correction methods make other computer vision tasks a multi-step approach and lack performance. In this work, we propose mitigating perspective distortion (MPD) by employing a fine-grained parameter control on a specific family of Möbius transform to model real-world distortion without estimating camera intrinsic and extrinsic parameters and without the need for actual distorted data. Also, we present a dedicated perspective distorted benchmark dataset, ImageNet-PD, to benchmark the robustness of deep learning models against this new dataset. The proposed method outperforms existing benchmarks, ImageNet-E and ImageNet-X. Additionally, it significantly improves performance on ImageNet-PD while consistently performing on standard data distribution. Notably, our method shows improved performance on three PD-affected real-world applications—crowd counting, fisheye image recognition, and person re-identification—and one PD-affected challenging CV task: object detection. The source code, dataset, and models are available on the project webpage at <https://prakashchhipa.github.io/projects/mpd>.

Keywords: Perspective Distortion · Self-supervised Learning · Robust Representation Learning

1 Introduction

Perspective distortion (PD) in real-world imagery is omnipresent and poses challenges in developing computer vision applications. PD arises from various factors, including camera positioning, depth, intrinsic parameters such as focal length and lens distortion, and extrinsic parameters such as rotation and translation. These factors collectively influence the projection of 3D scenes onto 2D planes [42], altering semantic interpretation



Fig. 1: MPD synthesizing perspective distortion with different orientations corresponding to parameter $c : (c_{real}, c_{imag})$.

and local geometry. Precisely estimating these parameters for perspective distortion correction is challenging, presenting a critical barrier to developing robust computer vision (CV) methods.

Previous studies in distortion correction predominantly concentrated on the correction of PD, with limited emphasis on enhancing the robustness of CV applications. Consequently, these correction methods transform CV tasks, such as image recognition and scene understanding, into a two-stage process: firstly, rectifying the distortion, and secondly, engaging in task-specific learning. GeoNet [29] employs CNN to predict distortion flow in images without prior distortion-type knowledge. PCL [57] uses perspective crop layers in CNNs to correct perspective distortions for 3D pose estimation. A cascaded deep structure network [46] corrects wide-angle portrait distortions without calibrated camera parameters. Another study [62] predicts per-pixel displacement for face portrait undistortion. PerspectiveNet [25] and ParamNet [25] predict perspective fields and derive camera calibration parameters, respectively. Methods for fisheye image distortion correction are also developed [54, 56]. Zolly [51] corrects perspective distortion by adjusting focal lengths in human mesh reconstruction tasks.

Since the correction methods are fundamentally not focused on robust representation learning for CV tasks, they essentially follow an inefficient two-step process. One possible approach is to focus on developing task-specific methods where PD is noticeable, such as crowd counting [22, 55], object re-identification [2], autonomous vehicle [52], and recognizing texts in scene [39]. Further, [9] meta-learning based approach to adapt 3D human pose estimation to varying degrees of camera distortion. SPEC [26] tackled PD by estimating the camera parameters directly from in-the-wild images to improve the accuracy of 3D human pose predictions. These task-specific approaches fall short in providing generalizable and robust representation learning for PD mitigation, a crucial step for advancing CV research towards distortion robustness.

The core reason behind the lack of generalizable and robust representation learning CV methods to mitigate PD is the absence of a real-world PD-affected training dataset. Both capturing real-world PD data and synthesizing PD-affected image data using parameter estimation are far from reality. Although existing benchmarks (ImageNet-X [23], ImageNet-E [28]) are partially affected by PD, a dedicated benchmark is needed. Therefore, there is a need to artificially synthesize PD in existing data without estimating camera intrinsic and extrinsic parameters. Since PD is non-linear in nature, the focus should be shifted from affine to potential non-linear transformation.

One potential non-linear family of transformation is Möbius transform [1, 37] in complex theory. Due to its conformal nature (preserving angles during transformation), [63] utilized the Möbius transform for eight specific transformations, including various twists and spreads, as a data augmentation strategy to enhance model performance on standard data distributions. In other works, Möbius convolution [36], the authors design a Möbius-equivalent convolution operator in spherical CNN for shape

classification and image segmentation tasks. OmniZoomer [5] integrates Möbius transformations into a neural network, enabling movement and variable zoom on omnidirectional images while addressing the challenges of blur and aliasing associated with such transformations. MöbiusGCN [3] addresses the limitation of existing graph convolution network GCNs in explicitly encoding transformations between joints in 3D human pose estimation. None of these works focus on perspective distortion.

Due to the non-linear characteristics of Möbius transforming geometry in complex space while ensuring conformality, it is a potential approach to imitate perspective distortion when adapted to image data. Therefore, in the proposed mitigating perspective distortion (MPD) method, we mathematically model PD in the images by employing fine-grained parameter control on Möbius. It is governed by four parameters a, b, c and d (Eq. 1). MPD focuses on a specific family of transforms corresponding to the parameter c to synthesize PD artificially. The core contribution of our method, MPD, is controlled parametrization for orientation (Fig. 1) and varying intensity (Fig. 2), which is capable of representing real-world perspective distortion.

MPD transforms input image coordinates into complex vectors and performs a PD-specific family of Möbius transforms. Transformed complex vectors are then remapped to real-valued pixels, later discretized to obtain images mimicking perspective distortion while maintaining the integrity of visual concepts. Optionally, padding is applied to replace the black background with boundary pixel values to achieve another variant called integrated padding background variant (refer Fig. 15 in suppl. material). Notably, MPD achieves this without estimating intrinsic and extrinsic parameters, thus bypassing the need for actual distorted data.

Incorporating PD in computer vision is beneficial for robustness since it can be used as an augmentation, resulting in improved performance on different computer vision tasks. To address the non-availability of a suitable benchmark, we have developed a dedicated perspectively distorted benchmark dataset, **ImageNet-PD**, derived from the ImageNet [11] validation set to evaluate the robustness of deep learning models against perspective distortion. We showcase a lack of robustness of standard deep learning models against perspective distortion by benchmarking them on ImageNet-PD.

We employ ImageNet-PD in our proposed MPD method and extensively investigate it across supervised and self-supervised learning methods (SimCLR [7], DINO [6]). MPD demonstrates improved performance on existing benchmarks, ImageNet-E [28], and ImageNet-X [23]. MPD improves 10% on ImageNet-PD while consistently maintaining the performance on the original ImageNet validation set. MPD shows performance improvements when adapted to multiple real-world CV applications. Specifically, MPD advances crowd counting with a novel method, *MPD-AutoCrowd*, achieving the highest performance (to the best of our knowledge) with a mean absolute error of 50.81 on ShanghaiTech-Part-A [60] and 96.80 on UCFCC50 [21]. It demonstrates effective knowledge transfer in fish-eye image recognition, improving performance by 3% on the VOC-360 [14] dataset. MPD combined with a transformer-based self-supervised method, Clip-ReIdent [15], excels in person re-identification, achieving mean average precision of 97.02 and 98.30 with and without re-ranking, respectively, on the DeepSportradar dataset [48]. Also, MPD adapted for object detection with

FasterRCNN [43] and FCOS [47] methods gain up to 3% performance improvement on COCO [30] dataset. The main contributions of this paper are as follows:

1. We mathematically model perspective distortion within deep learning by developing MPD, a method that successfully mimics perspective distortion without estimating the intrinsic and extrinsic camera parameters.
2. We develop a dedicated perspective distorted dataset (ImageNet-PD) derived from the ImageNet [11] validation set to benchmark the robustness of computer vision models against PD to address the non-availability of a suitable benchmark.
3. We demonstrate MPD’s generalizability, transferability, and adaptability across different learning paradigms, architectures, and real-world applications.
4. Additionally, we develop the MPD-AutoCrowd method in crowd counting to augment crowd and corresponding ground truth in scene images.

2 Methodology

Möbius from complex theory [1, 37] is a family of non-linear and conformal transformations. It is defined on complex number z where transformation is governed by four parameters, namely a , b , c , and d . Möbius transformation, $\Phi(z)$ is given in Eq.1:

$$\Phi(z) = \frac{az + b}{cz + d}, \quad ad - bc \neq 0 \quad (1)$$

Non-linearity: Given Möbius transformations for variables z_1 and z_2 , Möbius transformation of their sum $z_1 + z_2$:

$$\Phi(z_1 + z_2) = \frac{a(z_1 + z_2) + b}{c(z_1 + z_2) + d} = \frac{az_1 + az_2 + b}{cz_1 + cz_2 + d}, \rightarrow \Phi(z_1 + z_2) \neq \Phi(z_1) + \Phi(z_2) \quad (2)$$

However, Eq. (2) is not equal to the sum of the individual transformations ($\Phi(z_1) + \Phi(z_2)$); therefore, *non-linear*.

Conformality: Conformal transformations preserve angles between curves. The angle between two curves at a point is given by the *argument of the complex ratio* [19] of their derivatives. Let $f(z)$ and $g(z)$ represent curves in complex space. Their Möbius transforms are $\Phi(f(z))$, and $\Phi(g(z))$. If transformation preserves this ratio, it preserves angles. Eq. 3 shows that the angles are preserved, hence conformal [37]. Theoretical details and theorem are in suppl. material (refer sec. 5 and 6).

$$\arg((f(z))/g(z)) = \arg\left(\frac{d}{dz}(\Phi(f(z))) / \frac{d}{dz}(\Phi(g(z)))\right) \quad (3)$$

Möbius transform has been used in [63], which defined parameter estimation for eight specific transformed appearances such as clockwise twist, clockwise half-twist, spread, spread twist, counterclockwise twist, counterclockwise half-twist, inverse, and inverse spread. These appearances are used as data augmentation in image classification on CIFAR and Tiny ImageNet datasets. However, their parameter estimation does not focus on exploring a specific family of Möbius transform to synthesize perspective distortion due to a lack of fine-grained parameter control.

2.1 Mitigating perspective distortion (MPD)

The proposed MPD method mathematically models perspective distortion in the image using a family of Möbius transform. First, it maps the discrete coordinates (x, y) of the input image I from pixel space to complex vector $z = x + iy$ in complex space. Second, it applies Möbius transform to z and achieves corresponding non-linear transformed representation $z_m = x_m + iy_m$ by notably controlling real and imaginary parts of parameter c in Eq. 1. Next, the transformed complex vector z_m is mapped to real-valued pixel space (x_m, y_m) , which is discretized to get discrete pixel coordinates (x_d, y_d) . Obtained discrete pixel coordinates provide transformed image I_{MPD} mimicking perspective distortion. Optionally, padding is applied to replace black background by boundary pixel to obtain integrated padding variant of transformed image. In this manner, MPD can synthesize perspective distortion artificially without estimating intrinsic and extrinsic parameters of perspective distortion. Noticeably, It eliminates the need for perspectively distorted real image data. The core part of MPD in modeling

Algorithm 1 MPD Method for Perspective Distortion

Input: Input image I , Möbius parameters c , fixed-value complex numbers of (a, b, d)



Output: Transformed image I_{MPD} representing perspective distortion

```

1: Initialize  $c \leftarrow c_{real} + i \cdot c_{imag}$  ▷ Set complex parameter  $c$ 
2: for each pixel coordinate  $(x, y)$  in  $I$  do
3:    $z \leftarrow x + iy$  ▷ Map pixel coordinates to complex vector
4:    $z_m \leftarrow \frac{(a \cdot z + b)}{(c \cdot z + d)}$  ▷ Apply Möbius transformation
5:    $x_m + iy_m \leftarrow z_m$  ▷ Extract real and imaginary parts as transformed coordinates
6:    $(x_d, y_d) \leftarrow \text{round}(x_m), \text{round}(y_m)$  ▷ Discretize to nearest pixel values
7:    $I_{MPD}(x_d, y_d) \leftarrow I(x, y)$  ▷ Assign pixel value to transformed coordinates
8: end for
9:  $I_{MPD\_IB} \leftarrow \text{Padding}(I_{MPD})$  ▷ optional step for integrated padding background variant
10: return  $I_{MPD}, I_{MPD\_IB}$ 

```



11:

the PD lies in controlling the transformation through parameter c . Parameter c , being a complex number has real (c_{real}) and imaginary (c_{imag}) components; it is noteworthy that these components introduce real-world varieties of PD in the input image. c_{real} perspectively distorts the in input image horizontally whereas c_{imag} perspectively distorts the in input image vertically. The sign of c_{real} dictates the direction ($+ve$ for left view and $-ve$ for right view) of distortion, whereas magnitude defines the intensity of applied distortion. Similarly, the sign of c_{imag} dictates the direction ($+ve$ for top view and $-ve$ for bottom view) of distortion, whereas magnitude defines the intensity of applied distortion. The impact of signs of c_{real} and imaginary c_{imag} is depicted in Fig. 1. The synthesis of the increasingly distorted left view of the input image is demonstrated in Fig. 2 by controlling the intensity of c_{real} component of MPD. Demonstrations of increasingly distorted right, top, bottom, and other views are provided in Fig. 17 in the suppl. material. The main contribution of our MPD method is its ability to control pa-

parameterization for varying intensity and orientation, effectively representing real-world perspective distortions. One of the ways to incorporate MPD is augmentation, which is investigated in this work. MPD is effortlessly employable to incorporate into the learning process to make networks robust against perspective distortion, materializing the proposed approach to shifting the focus from exclusion to inclusion of PD. The above-described method MPD is presented in Algorithm 1. MPD returned a transformed image of a 'cat' as shown in step 11. Four regions are marked with colors (green, red, blue, and yellow) to observe the correspondence between input and transformed images. Although the current work does not focus on other parameters (a , b , and d), their impact on transformation is worth mentioning. The real component of parameters a and d controls scaling, and the imaginary components control rotations. The real component of parameter b controls horizontal translations, and the imaginary component controls vertical translations. Their visualizations are provided in Fig. 18 in the suppl. material. The transformation parameters are set as: $a = 1.0 + 0.0i$, $b = 0.0 + 0.0i$, $d = 1.0 + 0.0i$. The real and imaginary components of c range from 0.2 to 0.3 to mimic perspective distortion. These values are uniformly applied in all experiments unless otherwise stated. The influence of MPD to control perspective distortion is carried out via parameter c and given probability P to which MPD is applied during training. We establish the



Fig. 2: Demonstrate controlled scaling of distortion in MPD-transformed image. distortion is controlled by the intensity of the real component of complex parameter c_{real} ranging from -0.1 to -0.5 to synthesize perspectively distorted left-views on example image of 'cat'.

following objectives to thoroughly assess the proposed methodology in the context of perspective distortion.

- **Objective 1- Robustness evaluation of existing models:** Investigation of existing trained models for robustness against perspective distortion by evaluating them on the PD-dedicated newly developed benchmark ImageNet-PD.
- **Objective 2- MPD's effects on supervised learning:** Investigation of the effect of MPD in making supervised learning approach robust against perspective distortion.
- **Objective 3- MPD's effects on self-supervised learning:** Investigation of the effect of MPD in learning robust self-supervised representations against PD.
- **Objective 4- MPD's generalizability:** Investigating the effectiveness of MPD on cross-domain adaptability on diverse PD-affected real-world CV applications.

3 Experiments & Results

Naming convention and experimental details: *MPD* defines the proposed transformation (shown in Fig. 1). In addition, we also have another variant, *MPD IB (integrated padding background)*, which replaces the black background with boundary foreground pixel values of the transformed image (shown in Fig. 16 in suppl. material). **ImageNet-PD** benchmark dataset has eight subsets (Fig. 3); four corresponding to four orientations

(left, right, top, bottom) with black background (PD-L, PD-R, PD-T, PD-B). The other four subsets have the same orientations with integrated padding background, namely PD-LI (PD Left -Integrated padding background), PD-RI, PD-TI, and PD-BI.

Standard ResNet50 denotes ResNet50 model [16] with ImageNet-1k weights from torchvision library [40]. *Supervised:MPD* and *Supervised:MPD IB* indicate ResNet50 models trained on ImageNet [11] with MPD and MPD IB as augmentation respectively with defined probability P , following training protocol in torchvision guide [40]. *ssl:MPD* and *ssl:MPD IB* refer to ResNet50 encoders trained on ImageNet [11] incorporating MPD and MPD IB additional augmentation in a self-supervised approach using contrastive learning method, SimCLR [7] with $P = 0.8$, batch size of 512, and linear learning rate. After self-supervised pretraining (on a relatively smaller batch size of 512), models are fine-tuned for classification following the torchvision [40] protocol. *ssl:DINO-MPD* refer to ViT small encoder trained on ImageNet [11] incorporating MPD transformation in a self-supervised approach using DINO [6] method, with $P = 0.8$ and batch size of 512, following linear evaluation protocol described in [6]. Details on MPD adaptations in CV applications are described in respective sections.

3.1 Existing models' robustness evaluation

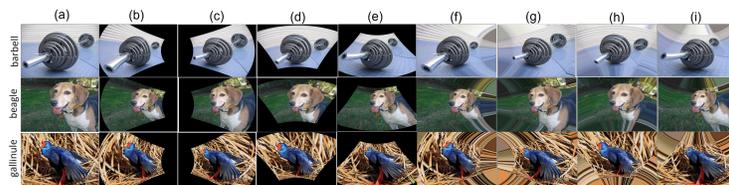


Fig. 3: Perspectively distorted image examples from ImageNet-PD benchmark dataset. (a) Original image, (b) Left view (PD-L), (c) Right view (PD-R), (d) Top view (PD-T), (e) Bottom view (PD-B), (f) Left view with integrated padding background (PD-LI), (g) Right view with integrated padding background (PD-RI), (h) Top view with integrated padding background (PD-TI), (i) Bottom view with integrated padding background (PD-BI).

New benchmark dataset- ImageNet-PD: To introduce perspective distortion by mimicking through the proposed transform MPD and its variant MPD-IB, we developed a new benchmarking dataset, ImageNet-PD (Fig. 3), derived from the ImageNet validation set. ImageNet-PD has eight subsets (Fig. 3), four corresponding to four orientations (left, right, top, bottom) with black background (PD-L, PD-R, PD-T, PD-B). The other four subsets have the same orientations but integrated padding backgrounds using boundary pixels (PD-LI, PD-RI, PD-TI, PD-BI). In our evaluation across various deep learning architectures, from classical CNNs like ResNet, VGG, and AlexNet to modern structures like EfficientNet and Vision Transformers (ViT), we consistently observe a significant drop in model accuracy when faced with images exhibiting perspective distortions, refer Fig. 4 (also in Tables 7, 8 in suppl. material). AlexNet shows a drop from 56.52% to 37.25% in Top-1 and from 79.06% to 61.04% in Top-5; EfficientNet-b2 falls from 80.60% to 66.71% in Top-1 and from 95.31% to 86.30% in Top-5; and even Vision Transformer models like ViT-B-16 decline from 81.07% to 74.09% in Top-1 and

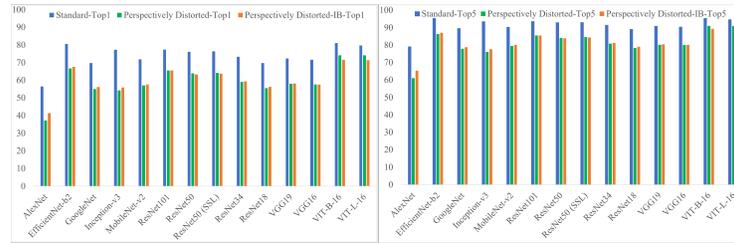


Fig. 4: Top1 and Top5 accuracies of ImageNet trained models with standard architectures on ImageNet-PD subsets. Blue bars shows performance on original ImageNet validation set, green bars shows mean performance on ImageNet-PD subsets with black background, and orange bars shows mean performance on ImageNet-PD subsets with integrated padding background.

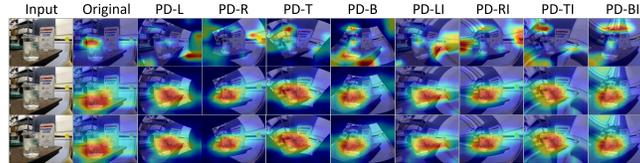


Fig. 5: Activation maps of the 'beaker' example in ImageNet-PD subsets. Standard ResNet50 model (row 1), supervised:MPD (row 2), and ssl:MPD (row 3).

from 95.32% to 90.98% in Top-5, as demonstrated by activation maps for a standard ResNet50 (Fig. 5, row 1). The uniform decrease in Top-1 and Top-5 accuracies across various models on ImageNet-PD vs. original ImageNet shows perspective distortion as a real challenge in computer vision.

3.2 MPD's effects on supervised learning

Existing Benchmarks: Both *supervised:MPD* and *supervised:MPD IB* models are trained on ImageNet dataset and evaluated on existing benchmarks and ImageNet-PD. ImageNet-X [23] and ImageNet-E [28]) considered as they are partially affected by perspective distortion. **ImageNet-E** [28] is a recent benchmark dataset specifically designed to evaluate the robustness of image classifiers concerning object attributes, including background settings, object sizes, spatial positioning, and orientation. **ImageNet-X** [23] introduces a set of sixteen human annotations focusing on attributes like pose, background, and lighting derived from the ImageNet. Fig. 6 demonstrates that both models improved performance on ImageNet-E by and reduced the accuracy drop (improving absolute accuracy) in *position* (2.91%), *direction* (2.14%), and *size* (1.19%) changes; supported by activation maps in Fig. 7. Other variants of MPD in supervised and self-supervised approaches consistently follow the robustness trend. Results are in Table 11 and Fig. 21 in suppl. material. Both *supervised:MPD* and *supervised:MPD IB* models outperform standard ResNet50 on the ImageNet-X dataset [23], achieving 63.23% accuracy compared to ResNet50's 60.64% and a lower error ratio of 1.44 versus 1.55 (refer Fig. 8). This improvement is particularly evident in handling perspective-related factors like size and shape variations. Detailed results are in Table 12 in suppl. material.



Fig. 6: ImageNet-E [28]: Drop of Top1 accuracy under background changes, size changes, random position (rp), random direction (rd), and average over 11 subsets. Lower is better. The mean is reported for size and background-related subsets. *Average* reports the mean of all subsets.

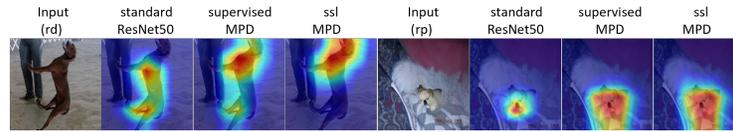


Fig. 7: Activation maps (ImageNet-E). 'redbone' from rd and 'peke' from rp subset.



Fig. 8: ImageNet-X [23]: Error ratio (close to 1.0 is best) of 16 factors in the plot. Accuracy is reported in numbers for each subset. *Average* reports the mean accuracy of all factors. Both supervised:MPD and ssl:MPD models improved the error ratio with higher accuracy.

Benchmarking on ImageNet-PD: Both *supervised:MPD* and *supervised:MPD IB* demonstrate notable improvements (Fig. 9 (a)-(b)) on ImageNet-PD subsets, achieving average Top-1 accuracy enhancements of 9.42% and 9.54% respectively, over the standard ResNet50 model. Please refer to Table 1 and Table 10 for detailed comparison; Table 10 is in suppl. material. Concurrently, these models maintain performance on the original ImageNet validation set, highlighting their efficacy and adaptability amidst perspective distortion. The activation maps shown in Fig. 5 (row 2) further corroborate this outperforming trend. Ablation on the probability of applying MPD and MPD IB as a data augmentation (Fig. 9 (c)-(d)) exhibits consistent performance. MPD also compared with other popular augmentation methods (Mixup [58], Cutout [12], AugMix [17], and Pixmix [18]) on ImageNet-PD and ImageNet-E, consistently outperforming them. Detailed results are in the suppl. material (refer to Table 17).

Table 1: Comparative analysis of MPD trained on supervised (*supervised:MPD model*) and self-supervised (*ssl:MPD model*) approaches and evaluated on original ImageNet and **ImageNet-PD subsets** (black background). The probability P of applying MPD in model fine-tuning is in the first column. More results on probabilities 0.0 and 1.0 in Table 9 in suppl. material.

P	Original Validation Set		Perspectively Distorted Top-view (PD-T)		Perspectively Distorted Bottom-view (PD-B)		Perspectively Distorted Left-view (PD-L)		Perspectively Distorted Right-view (PD-R)	
	Top1	Top 5	Top1	Top 5	Top1	Top 5	Top1	Top 5	Top1	Top 5
Supervised training from scratch										
0.0	76.13±0.04	92.86±0.01	63.37±0.06	83.61±0.02	61.15±0.04	81.86±0.01	65.20±0.03	85.13±0.03	65.84±0.06	85.64±0.02
0.2	76.05±0.03	92.99±0.02	72.48±0.02	91.02±0.01	72.12±0.02	91.08±0.02	72.57±0.03	91.13±0.03	72.80±0.05	91.35±0.02
0.4	76.17±0.04	93.03±0.02	73.13±0.03	91.33±0.01	72.94±0.02	91.42±0.01	73.19±0.06	91.48±0.01	73.38±0.01	91.69±0.01
0.6	76.19±0.05	93.14±0.03	73.23±0.02	91.46±0.03	73.01±0.05	91.34±0.02	73.47±0.02	91.66±0.01	73.54±0.06	91.61±0.02
0.8	76.34±0.02	93.03±0.02	73.00±0.02	90.69±0.02	72.31±0.03	90.81±0.03	73.50±0.04	91.33±0.03	72.91±0.02	91.29±0.02
Self-supervised pre-taining on contrastive learning integrating MPD (probability for pre-training=0.8)										
0.2	76.37±0.05	93.60±0.02	72.34±0.02	90.81±0.02	72.24±0.03	90.95±0.02	72.48±0.04	91.08±0.01	72.81±0.03	91.16±0.01
0.4	76.77±0.02	93.40±0.01	73.23±0.03	91.39±0.01	73.06±0.04	91.54±0.02	73.55±0.03	91.53±0.03	73.54±0.05	91.65±0.03
0.6	76.14±0.04	93.58±0.02	73.39±0.02	91.50±0.02	73.29±0.03	91.57±0.01	73.57±0.03	91.66±0.02	73.73±0.04	91.58±0.02
0.8	76.29±0.03	92.78±0.02	73.61±0.04	91.49±0.03	73.27±0.05	91.48±0.03	73.60±0.06	91.69±0.01	73.81±0.02	91.71±0.02

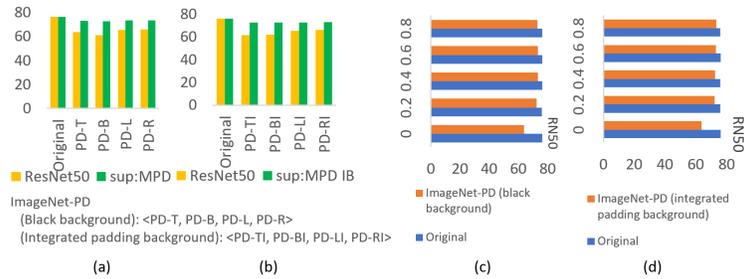


Fig. 9: (a)-(b)[x-axis: ImageNet-PD subsets, y-axis: accuracy] → Top-1 accuracy comparison of standard ResNet50 & *supervised:MPD* and *supervised:MPD IB* models on ImageNet-PD subsets. (c)-(d)[x-axis: accuracy, y-axis: probability] → Mean top-1 accuracy across ImageNet-PD subsets for *supervised:MPD* and *supervised:MPD IB* models with varied probability. 0.0 refers standard ResNet50. Original-original ImageNet validation set.

3.3 MPD’s effects on self-supervised learning

MPD as data augmentation integrated in SimCLR [7] and DINO [6] (small ViT [13] backbone), as in Fig. 19 and 20 in suppl. material. Models are pre-trained and finetuned / liner-probed on ImageNet and evaluated on existing benchmarks and ImageNet-PD.

Existing Benchmarks: *ssl:MPD* and *ssl:MPD IB* models demonstrate efficacy in handling perspective-related variations in **ImageNet-E** as shown in Fig. 6. Specifically, it shows a reduction in the accuracy drop for position (1.21%), direction (1.41%), and size (0.48%) changes compared to the standard ResNet50, indicating its robustness. On the **ImageNet-X** benchmark, *ssl:MPD* and *ssl:MPD IB* models outperform with an average accuracy of 62.99%, compared to the standard ResNet50’s 60.64%, and also show a reduced error ratio of 1.41 against the ResNet50’s 1.55 (refer Fig. 8). Notably, performance improvement is evident in managing variations in size and shape, showcasing its robustness against PD. Detailed results are in Table 12 in suppl. material.

Benchmarking on ImageNet-PD: *ssl:MPD* model (Fig. 10 (a)) shows an average improvement of approximately 10.32% across the PD subsets, while the *ssl:MPD IB*

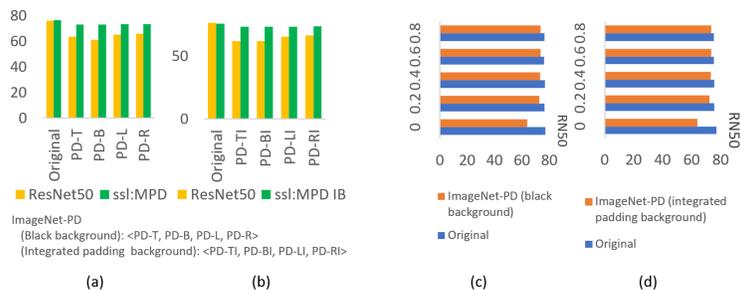


Fig. 10: (a)-(b)[x-axis: ImageNet-PD subsets, y-axis: accuracy] \rightarrow Top-1 accuracy comparison of standard ResNet50 and *ssl:MPD* & *ssl:MPD IB* models on ImageNet-PD subsets. (c)-(d)[x-axis: accuracy, y-axis: probability] \rightarrow Mean top-1 accuracy across ImageNet-PD subsets for *ssl:MPD* model with varied probability. 0.0 refers standard ResNet50.

(Fig. 10 (b)) demonstrates an average increase of around 10.14%. In addition, both variants display robust performance on the original ImageNet validation set. Please refer to Table 1 and Table 10 for detailed comparison in suppl. material. The activation maps depicted in Fig. 5 (row 3) further validate the improved performance. Moreover, ablation on the probability of applying MPD and MPD IB as data augmentation (Fig. 10 (c)-(d)) demonstrates a consistent increase in performance. We have also performed linear evaluations on self-supervised method SimCLR [7] and DINO [6] (refer Table 15, 16; Fig. 22 for detailed results in suppl. material). Similarly, significant improvement is noticed on ImageNet-PD subsets while maintaining the performance on standard ImageNet as shown in Fig. 11 (refer Table 13, 14 for detailed results in suppl. material).

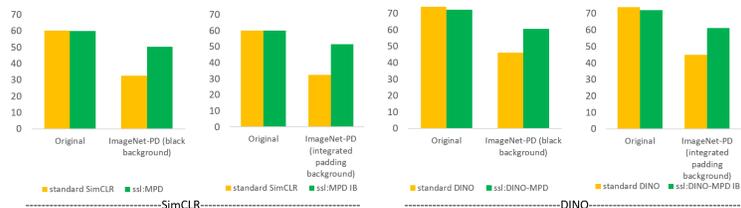


Fig. 11: Linear evaluation of SimCLR [7] and DINO [6] with MPD and MPD IB augmentations.

3.4 MPD's generalizability

This section discusses MPD's adaptability in diverse applications, namely crowd counting (CC), fisheye, and person re-identification, and challenging CV tasks such as object detection, where perspective distortion is evident.

Crowd counting: Numerous human heads within a single image present challenges due to the variations in size, orientation, and shape, naturally including perspective distortion. We investigate the potential of MPD in crowd counting on several publicly available datasets, including ShanghaiTech (SHTech) Part A and Part B [60] and UCF-CC50 [21]. The ShanghaiTech Part A features highly congested scenes, and Part B

consists of sparser scenes captured on a busy street. The UCF-CC50 dataset comprises 50 diverse scenes featuring a considerable variation in the crowd.

MPD-CC: MPD is adapted to CC by extending its capability to transform scene images and their CC labels, as mentioned in steps 4-7 of algorithm 2.

MPD-AutoCrowd: We introduce MPD-AutoCrowd (Algorithm 2), a novel approach for the automated augmentation of the crowd in scene images by upgrading MPD-CC. It seamlessly integrates synthetic crowds into the background regions of transformed images, complete with corresponding labels. This approach facilitates scene-level crowd augmentation without human supervision (Fig. 12). MPD-CC and MPD-AutoCrowd



Fig. 12: MPD-AutoCrowd: Artificially augments crowd & labels. (a) Original image with labels, (b) transformed image & labels, (c) Transformed image with augmented crowd, (d) Transformed image with augmented crowd & labels (yellow box).

are integrated with P2P-Net [45] to evaluate the effectiveness of MPD in crowd counting. Training parameters, CNN architecture (VGG16), and evaluation configuration are used as mentioned in P2P-Net [45]. The probability of applying MPD-CC and MPD AutoCrowd is 0.1, and real and imaginary components of c range from 0.2 to 0.6. This method has yielded improved results, surpassing previous state-of-the-art performance. Refer to Table 2 for a detailed comparison with other CC methods. MPD-CC and MPD-AutoCrowd outperform on multiple benchmarks. MPD-AutoCrowd has a

Algorithm 2 MPD-AutoCrowd

Input: Image I , CC Points P_o , Möbius parameters c , fixed-value complex numbers of (a, b, d)
Output: Transformed image I' , Transformed crowd counting points P_o'

1: $I' \leftarrow \text{MPD}(I, a, b, c, d)$		▷ call MPD algorithm
2: $P_o' \leftarrow []$		▷ initialize points list
3: for each point (x, y) in P_o do		
4: $z_{point} \leftarrow x + i \cdot y$		▷ point to complex number
5: $z'_{point} \leftarrow \frac{(a \cdot z_{point} + b)}{(c \cdot z_{point} + d)}$		▷ apply Möbius to point
6: $x', y' \leftarrow \Re(z'_{point}), \Im(z'_{point})$		▷ transformed point
7: $P' \leftarrow P_o' \cup \{(x', y')\}$		▷ transformed point
8: $x'_{ref}, y'_{ref} \leftarrow \text{Reflect}(x', y')$		▷ reflect point
9: $P_o' \leftarrow P_o' \cup \{(x'_{ref}, y'_{ref})\}$		▷ reflected points
10: end for		
11: $P_o'_{agg} \leftarrow \text{mean}(P_o')$		▷ aggregate points
12: return $I', P_o'_{agg}$		▷ transformed image and points

MAE of 50.81 on SHTech Part A and 6.73 on SHTech Part B. The performance on the UCF-CC50 (a highly dense crowds) dataset, where MPD-AutoCrowd has an MAE of 96.80 and a Mean Squared Error (MSE) of 139.50, significantly outperforms other models. We further improved results on UCF-CC50 with ResNet50 encoder as shown in Table 18 in suppl. material. We also adapted other perspective transforms [59] [38]

and RandAug [10] augmentation for crowd counting and compared with MPD-CC and MPD-AutoCrowd which shows that our proposed methods outperform. Details are in suppl. material (refer Table 19 and Fig. 23 in sec. 14.1).

Table 2: Comparison of different crowd-counting methods with the proposed work. MPD-CC & MPD-AutoCrowd outperform on multiple benchmarks

Method	Venue	SHTech PartA		SHTech PartB		UCF_CC-50	
		MAE	MSE	MAE	MSE	MAE	MSE
CAN [31]	CVPR'19	62.3	100	7.8	12.2	212.2	243.7
Bayesian+ [34]	ICCV'19	62.8	101.8	7.7	12.7	229.3	308.2
S-DCNet [53]	ICCV'19	58.3	95	6.7	10.7	204.2	301.3
SA+SPANet [8]	ICCV'19	59.4	100	6.5	9.9	232.6	311.7
SDANet [35]	AAAI'20	63.6	101.8	7.8	10.2	227.6	316.4
ADSCNet [4]	CVPR'20	55.4	97.7	6.4	11.3	198.4	267.3
ASNet [24]	CVPR'20	57.78	90.13	-	-	174.84	251.63
AMRNet [32]	ECCV'20	61.59	98.36	7.02	11	184	265.8
AMNet [20]	ECCV'20	56.7	93.4	6.7	10.2	208.4	297.3
DM [49]	NeurIPS'20	59.7	95.7	7.4	11.8	211	291.5
P2P-Net [45]	ICCV'21	52.74	85.06	6.25	9.9	172.72	256.18
SDA(best) [33]	ICCV'21	52.90	87.30	-	-	159.1	239.4
ChfL [44]	CVPR'22	57.50	94.30	6.90	11.0	-	-
DMCNet [50]	WACV'23	58.46	84.55	8.64	13.67	-	-
MPD-CC	ours	51.93	84.3	6.73	9.82	101.30	140.65
MPD-AutoCrowd	ours	50.81	85.01	6.61	9.58	96.80	139.50

Transfer learning on fisheye images Fisheye image distortions, caused by a wide field of view, are handled by spherical perspective models in [27, 61]. MPD simulates spherical views in pixel space and improves representation learning for fisheye images. We showcase the transferable capability of models trained with MPD and MPD IB through supervised and self-supervised methods on the ImageNet dataset, applying them to the VOC-360 fisheye dataset [14]. This dataset has 39,575 fisheye images with multi-label classification labels. We fine-tune these models on the VOC360, running for 100 epochs with a batch size 16. We show the performance of supervised and self-supervised models for the multi-label classification task in Table 3. ResNet50 is used as backbone. Ablations on label efficiency (Table 20, 21, and Fig. 24) are in suppl. material. *supervised:MPD* and *supervised:MPD IB* models fine-tuned on VOC-360 dataset

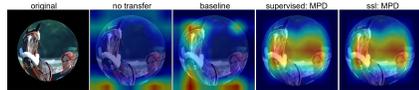


Fig. 13: Activation maps of 'horse' and 'person' on VOC-360 dataset.

Table 3: MPD performance on VOC-360.

Transfer Learning	Accuracy	Precision	Recall
No transfer learning	82.02	94.90	89.96
Baseline	91.32	97.24	95.67
supervised: MPD IB	91.87	97.00	96.22
supervised: MPD	92.69	97.76	96.23
ssl: MPD IB	92.43	97.16	96.50
ssl: MPD	94.79	98.17	97.44

achieve accuracies 92.69% and 91.87% respectively. *ssl:MPD* and *ssl:MPD IB* models fine-tuned on VOC-360 achieve 92.43% and 94.79% respectively. *ssl:MPD* performs best among all models. Fig. 13 depicts the activation maps.

Person Re-Identification: The DeepSportRadar re-identification dataset [48] includes video frames with varying poses and camera angles, offering training (436 queries,

8133 gallery images) and testing sets (50 queries, 910 gallery images) to address perspective variability in re-identification tasks. We incorporated the MPD within the state-of-the-art CLIP-ReIdent method [15]. The probability of applying MPD is 0.1, and real & imaginary components of c range from 0.2 to 0.6. CLIP-ReIdent is CLIP [41]-adapted contrastive image-image pre-training for person re-identification, leveraging transformer architecture. MPD-incorporated Clip-ReIdent models improve performance with both backbone architectures, ViT-L-14 and ResNet50x16 (ref. Table 4). Ablations on MPD parameters are in suppl. material in Table 22. In addition, it boosts performance across different frame counts, achieving competitive results with limited frames and performance further enhanced by longer training (refer Table 5).

Table 4: MPD improves the performance of the person re-identification. RR:re-ranking.

Method	Encoder	mAP (w/o RR)	mAP (with RR)
Baseline		72.70	-
Clip-ReIdent	ViT-L-14	96.90	98.20
MPD (Clip-ReIdent)		97.02	98.30
Clip-ReIdent	ResNet50x16	88.50	94.90
MPD (Clip-ReIdent)		91.95	97.50

Table 5: Label efficiency shown in terms of frame counts.

epoch		6 frames	10 frames	20 frames
8	w/o re-ranking	95	95.19	97.02
	with re-ranking	98	98.05	98.3
30	w/o re-ranking	95.68	96.42	97.02
	with re-ranking	98.24	98.27	98.3

Object Detection: We applied MPD method for the object detection (MPD-OD) by transforming image and bounding boxes as augmentation to mimic PD (fig. 14). We have trained FasterRCNN [43] and FCOS [47] with MPD-OD on COCO [30]. Applying MPD-OD to 50% probability ($P=0.5$), we achieved 3% and 1.6% gain over original models (Table 6). Ablations are in suppl. material (Table 23 and Fig. 25 in sec. 14.4).

Method	Original	MPD-OD (IoU=0.50:0.95/IoU=0.50)
FasterRCNN [43]	37.10/55.80	40.00/61.10
FCOS [47]	38.60/57.40	40.20/60.30

Table 6: MPD-OD on COCO dataset.

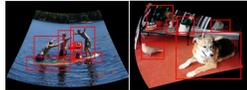


Fig. 14: MPD-OD examples

4 Conclusion

We address perspective distortion for robust computer vision applications. The proposed MPD method models real-world distortions by applying a family of Möbius transforms on existing data without estimating camera parameters or relying on actual distorted data. We develop a perspective-distorted ImageNet-PD dataset to benchmark the robustness of the computer vision models and reveal that existing models lack robustness. Further, MPD-incorporated models improve performance on existing benchmarks ImageNet-E and ImageNet-X and significantly enhance results on the ImageNet-PD while maintaining consistent performance on standard data. Also, MPD incorporated in practical applications such as crowd counting, fisheye image recognition, person re-identification, and object detection improved performance. We believe that our work on mitigating PD could inspire future work for other CV challenges.

Acknowledgment The authors thank Sumit Rakesh, Luleå University of Technology, for his support with the Lotty Bruzelius cluster. We also thank the National Supercomputer Centre at Linköping University for the Berzelius supercomputing, supported by the Knut and Alice Wallenberg Foundation.

References

1. Arnold, D.N., Rogness, J.P.: Möbius transformations revealed. *Notices of the American Mathematical Society* **55**(10), 1226–1231 (2008)
2. Ayala-Acevedo, A., Devgun, A., Zahir, S., Askary, S.: Vehicle re-identification: Pushing the limits of re-identification. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*. vol. 2 (2019)
3. Azizi, N., Possegger, H., Rodolà, E., Bischof, H.: 3d human pose estimation using möbius graph convolutional networks. In: *Proceedings of the European Conference on Computer Vision*. pp. 160–178. Springer (2022)
4. Bai, S., He, Z., Qiao, Y., Hu, H., Wu, W., Yan, J.: Adaptive dilated network with self-correction supervision for counting. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 4594–4603 (2020)
5. Cao, Z., Ai, H., Cao, Y.P., Shan, Y., Qie, X., Wang, L.: Omnizoomer: Learning to move and zoom in on sphere at high-resolution. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 12897–12907 (2023)
6. Caron, M., Touvron, H., Misra, I., Jégou, H., Mairal, J., Bojanowski, P., Joulin, A.: Emerging properties in self-supervised vision transformers. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 9650–9660 (2021)
7. Chen, T., Kornblith, S., Norouzi, M., Hinton, G.: A simple framework for contrastive learning of visual representations. In: *Proceedings of the International Conference on Machine Learning*. pp. 1597–1607. PMLR (2020)
8. Cheng, Z.Q., Li, J.X., Dai, Q., Wu, X., Hauptmann, A.G.: Learning spatial awareness to improve crowd counting. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 6152–6161 (2019)
9. Cho, H., Cho, Y., Yu, J., Kim, J.: Camera distortion-aware 3d human pose estimation in video with optimization-based meta-learning. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 11169–11178 (2021)
10. Cubuk, E.D., Zoph, B., Shlens, J., Le, Q.V.: Randaugment: Practical automated data augmentation with a reduced search space. In: *CVPR workshops*. pp. 702–703 (2020)
11. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: *2009 IEEE conference on computer vision and pattern recognition*. pp. 248–255. Ieee (2009)
12. DeVries, T., Taylor, G.W.: Improved regularization of convolutional neural networks with cutout. *arXiv preprint arXiv:1708.04552* (2017)
13. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al.: An image is worth 16x16 words: Transformers for image recognition at scale. In: *Proceedings of the International Conference on Learning Representations* (2020)
14. Fu, J., Bajić, I.V., Vaughan, R.G.: Datasets for face and object detection in fisheye images. *Data in brief* **27**, 104752 (2019)
15. Habel, K., Deuser, F., Oswald, N.: Clip-reident: Contrastive training for player re-identification. In: *Proceedings of the 5th International ACM Workshop on Multimedia Content Analysis in Sports*. pp. 129–135 (2022)
16. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 770–778 (2016)
17. Hendrycks, D., Mu, N., Cubuk, E.D., Zoph, B., Gilmer, J., Lakshminarayanan, B.: Augmix: A simple data processing method to improve robustness and uncertainty. In: *International Conference on Learning Representations*

18. Hendrycks, D., Zou, A., Mazeika, M., Tang, L., Li, B., Song, D., Steinhardt, J.: Pixmix: Dreamlike pictures comprehensively improve safety measures. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 16783–16792 (2022)
19. Howie, J.M.: Complex analysis. Springer Science & Business Media (2003)
20. Hu, Y., Jiang, X., Liu, X., Zhang, B., Han, J., Cao, X., Doermann, D.: Nas-count: Counting-by-density with neural architecture search. In: Proceedings of the European Conference on Computer Vision. pp. 747–766. Springer (2020)
21. Idrees, H., Saleemi, I., Seibert, C., Shah, M.: Multi-source multi-scale counting in extremely dense crowd images. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2547–2554 (2013)
22. Idrees, H., Tayyab, M., Athrey, K., Zhang, D., Al-Maadeed, S., Rajpoot, N., Shah, M.: Composition loss for counting, density map estimation and localization in dense crowds. In: Proceedings of the European Conference on Computer Vision. pp. 532–546 (2018)
23. Idrissi, B.Y., Bouchacourt, D., Balestrierio, R., Evtimov, I., Hazirbas, C., Ballas, N., Vincent, P., Drozdal, M., Lopez-Paz, D., Ibrahim, M.: Imagenet-x: Understanding model mistakes with factor of variation annotations. Proceedings of the International Conference on Learning Representations (2023)
24. Jiang, X., Zhang, L., Xu, M., Zhang, T., Lv, P., Zhou, B., Yang, X., Pang, Y.: Attention scaling for crowd counting. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4706–4715 (2020)
25. Jin, L., Zhang, J., Hold-Geoffroy, Y., Wang, O., Blackburn-Matzen, K., Sticha, M., Fouhey, D.F.: Perspective fields for single image camera calibration. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 17307–17316 (2023)
26. Kocabas, M., Huang, C.H.P., Tesch, J., Müller, L., Hilliges, O., Black, M.J.: Spec: Seeing people in the wild with an estimated camera. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 11035–11045 (2021)
27. Kumar, V.R., Eising, C., Witt, C., Yogamani, S.: Surround-view fisheye camera perception for automated driving: Overview, survey & challenges. IEEE Transactions on Intelligent Transportation Systems (2023)
28. Li, X., Chen, Y., Zhu, Y., Wang, S., Zhang, R., Xue, H.: Imagenet-e: Benchmarking neural network robustness via attribute editing. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 20371–20381 (2023)
29. Li, X., Zhang, B., Sander, P.V., Liao, J.: Blind geometric distortion correction on images through deep learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4855–4864 (2019)
30. Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft coco: Common objects in context. In: Computer Vision–ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part V 13. pp. 740–755. Springer (2014)
31. Liu, W., Salzmann, M., Fua, P.: Context-aware crowd counting. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5099–5108 (2019)
32. Liu, X., Yang, J., Ding, W., Wang, T., Wang, Z., Xiong, J.: Adaptive mixture regression network with local counting map for crowd counting. In: Proceedings of the European Conference on Computer Vision. pp. 241–257. Springer (2020)
33. Ma, Z., Hong, X., Wei, X., Qiu, Y., Gong, Y.: Towards a universal model for cross-dataset crowd counting. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 3205–3214 (2021)

34. Ma, Z., Wei, X., Hong, X., Gong, Y.: Bayesian loss for crowd count estimation with point supervision. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 6142–6151 (2019)
35. Miao, Y., Lin, Z., Ding, G., Han, J.: Shallow feature based dense attention network for crowd counting. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 34, pp. 11765–11772 (2020)
36. Mitchel, T.W., Aigerman, N., Kim, V.G., Kazhdan, M.: Möbius convolutions for spherical cnns. In: Proceedings of the ACM SIGGRAPH Conference. pp. 1–9 (2022)
37. Olsen, J.: The geometry of möbius transformations. Rochester: University of Rochester (2010)
38. Papakipos, Z., Bitton, J.: Augly: Data augmentations for adversarial robustness. In: CVPR. pp. 156–163 (2022)
39. Phan, T.Q., Shivakumara, P., Tian, S., Tan, C.L.: Recognizing text with perspective distortion in natural scenes. In: Proceedings of the IEEE/CVF international Conference on Computer Vision. pp. 569–576 (2013)
40. PyTorch: Vision: Datasets, transforms and models specific to computer vision (2023), <https://github.com/pytorch/vision/tree/main/references/classification>, original repository for PyTorch Vision Reference Scripts, ACCESSED: August, 1, 2023
41. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: Proceedings of the International Conference on Machine Learning. pp. 8748–8763. PMLR (2021)
42. Rahman, T., Krouglicof, N.: An efficient camera calibration technique offering robustness and accuracy over a wide range of lens distortion. *IEEE Transactions on Image Processing* **21**(2), 626–637 (2011)
43. Ren, S., He, K., Girshick, R., Sun, J.: Faster r-cnn: Towards real-time object detection with region proposal networks. *Advances in neural information processing systems* **28** (2015)
44. Shu, W., Wan, J., Tan, K.C., Kwong, S., Chan, A.B.: Crowd counting in the frequency domain. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 19618–19627 (2022)
45. Song, Q., Wang, C., Jiang, Z., Wang, Y., Tai, Y., Wang, C., Li, J., Huang, F., Wu, Y.: Rethinking counting and localization in crowds: A purely point-based framework. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 3365–3374 (2021)
46. Tan, J., Zhao, S., Xiong, P., Liu, J., Fan, H., Liu, S.: Practical wide-angle portraits correction with deep structured models. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3498–3506 (2021)
47. Tian, Z., Shen, C., Chen, H., He, T.: Fcos: Fully convolutional one-stage object detection. In: 2019 IEEE/CVF International Conference on Computer Vision (ICCV). pp. 9626–9635 (2019). <https://doi.org/10.1109/ICCV.2019.00972>
48. Van Zandycke, G., Somers, V., Istasse, M., Don, C.D., Zambrano, D.: Deepsporadar-v1: Computer vision dataset for sports understanding with high quality annotations. In: Proceedings of the 5th International ACM Workshop on Multimedia Content Analysis in Sports. pp. 1–8 (2022)
49. Wang, B., Liu, H., Samaras, D., Nguyen, M.H.: Distribution matching for crowd counting. *Proceedings of the Advances in Neural Information Processing Systems* **33**, 1595–1607 (2020)
50. Wang, M., Cai, H., Dai, Y., Gong, M.: Dynamic mixture of counter network for location-agnostic crowd counting. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 167–177 (2023)

51. Wang, W., Ge, Y., Mei, H., Cai, Z., Sun, Q., Wang, Y., Shen, C., Yang, L., Komura, T.: Zolly: Zoom focal length correctly for perspective-distorted human mesh reconstruction. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 3925–3935 (2023)
52. Wang, Y., Fathi, A., Kundu, A., Ross, D.A., Pantofaru, C., Funkhouser, T., Solomon, J.: Pillar-based object detection for autonomous driving. In: Proceedings of the European Conference on Computer Vision. pp. 18–34. Springer (2020)
53. Xiong, H., Lu, H., Liu, C., Liu, L., Cao, Z., Shen, C.: From open set to closed set: Counting objects by spatial divide-and-conquer. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 8362–8371 (2019)
54. Yang, S., Lin, C., Liao, K., Zhao, Y.: Innovating real fisheye image correction with dual diffusion architecture. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 12699–12708 (2023)
55. Yang, Y., Li, G., Wu, Z., Su, L., Huang, Q., Sebe, N.: Weakly-supervised crowd counting learns from sorting rather than locations. In: Proceedings of the European Conference on Computer Vision. pp. 1–17. Springer (2020)
56. Yin, X., Wang, X., Yu, J., Zhang, M., Fua, P., Tao, D.: Fisheyerecnet: A multi-context collaborative deep network for fisheye image rectification. In: Proceedings of the European Conference on Computer Vision. pp. 469–484 (2018)
57. Yu, F., Salzmann, M., Fua, P., Rhodin, H.: Pcls: Geometry-aware neural reconstruction of 3d pose with perspective crop layers. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 9064–9073 (2021)
58. Zhang, H., Cisse, M., Dauphin, Y.N., Lopez-Paz, D.: mixup: Beyond empirical risk minimization. In: International Conference on Learning Representations (2018)
59. Zhang, J., Zhang, J., Chen, B., Gao, J., Ji, S., Zhang, X., Wang, Z.: A perspective transformation method based on computer vision. In: ICAICA. pp. 765–768. IEEE (2020)
60. Zhang, Y., Zhou, D., Chen, S., Gao, S., Ma, Y.: Single-image crowd counting via multi-column convolutional neural network. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 589–597 (2016)
61. Zhang, Y., Song, J., Ding, Y., Yuan, Y., Wei, H.L.: FSD-BRIEF: A distorted BRIEF descriptor for fisheye image based on spherical perspective model. *Sensors (Basel)* (2021)
62. Zhao, Y., Huang, Z., Li, T., Chen, W., LeGendre, C., Ren, X., Shapiro, A., Li, H.: Learning perspective undistortion of portraits. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 7849–7859 (2019)
63. Zhou, S., Zhang, J., Jiang, H., Lundh, T., Ng, A.Y.: Data augmentation with mobius transformations. *Machine Learning: Science and Technology* **2**(2), 025016 (2021)

Möbius Transform for Mitigating Perspective Distortions in Representation Learning

Supplementary Material

5 Non-linearity and Conformality of Möbius

Non-linearity Given Möbius transformations for variables z_1 and z_2 , Möbius transformation of their sum $z_1 + z_2$:

$$\Phi(z_1) = \frac{az_1 + b}{cz_1 + d}, \quad \Phi(z_2) = \frac{az_2 + b}{cz_2 + d} \quad (4)$$

$$\Phi(z_1 + z_2) = \frac{a(z_1 + z_2) + b}{c(z_1 + z_2) + d} = \frac{az_1 + az_2 + b}{cz_1 + cz_2 + d} \quad (5)$$

$$\Phi(z_1 + z_2) \neq \Phi(z_1) + \Phi(z_2) \quad (6)$$

However, Eq. (6) is not equal to the sum of the individual transformations ($\Phi(z_1) + \Phi(z_2)$); therefore, *non-linear*.

Conformality Conformal transformations preserve angles between curves. The angle between two curves at a point is given by the *argument of the complex ratio* [19] of their derivatives. Let $f(z)$ and $g(z)$ represent curves in complex space. If transformation preserves this ratio, it preserves angles.

Möbius transformation Φ of $f(z)$ and $g(z)$ is given as:

$$\Phi(f(z)) = \frac{af(z) + b}{cf(z) + d}, \quad \Phi(g(z)) = \frac{ag(z) + b}{cg(z) + d} \quad (7)$$

To find the derivatives of the transformed functions, we apply the chain rule:

$$\begin{aligned} \Phi'(f(z)) &= \frac{d}{dz}(\Phi(f(z))) = \frac{d}{dz} \left(\frac{af(z) + b}{cf(z) + d} \right) \\ \Phi'(g(z)) &= \frac{d}{dz}(\Phi(g(z))) = \frac{d}{dz} \left(\frac{ag(z) + b}{cg(z) + d} \right) \end{aligned} \quad (8)$$

The angle between the original and transformed curves is given by the argument of the ratio [19]:

$$\arg((f(z))/g(z)) \equiv \arg \left(\frac{d}{dz}(f(z)) \right) - \arg \left(\frac{d}{dz}(g(z)) \right) \quad (9)$$

$$\begin{aligned} \arg(\Phi'(f(z))/\Phi'(g(z))) &\equiv \\ \arg \left(\frac{d}{dz}(\Phi(f(z))) \right) &- \arg \left(\frac{d}{dz}(\Phi(g(z))) \right) \end{aligned} \quad (10)$$

Eq. 11 indicates conformality [37].

$$\arg((f(z))/g(z)) = \arg(\Phi'(f(z))/\Phi'(g(z))) \quad (11)$$

6 Theorem: MPD complies semantic-preserving property in pixel space by approximating conformality

Proof: Consider $I : D \rightarrow \mathbb{C}$ as the discrete representation of an image over the domain D , corresponding to the pixel grid, and let $\Phi : \mathbb{C} \rightarrow \mathbb{C}$ be the continuous Möbius transformation given by $\Phi(z) = \frac{az+b}{cz+d}$. MPD applies discrete transformation T to each point p in D , which approximates Φ . The transformation T is constructed to preserve the angles between points in D as follows:

$$T(p) = \begin{cases} \frac{a \cdot I(p) + b}{c \cdot I(p) + d} & , c \cdot I(p) + d \neq 0, \end{cases} \quad (12)$$

The preservation of angles in the discrete domain is approximated by considering the discrete differential ΔI at a point p in D , defined as the vector of differences between $I(p)$ and $I(q)$ evaluated at neighboring points. For any two adjacent points $p, q \in D$, the discrete differential $\Delta I(p)$ and $\Delta I(q)$ form vectors that subtend an angle θ in the original pixel space. The transformation T aims to approximate the preservation of θ such that the angle between $\Delta T(p)$ and $\Delta T(q)$, denoted as $\hat{\theta}$, is close to θ in the transformed image space. This is expressed as $\Delta T(p)$ and $\Delta T(q)$ approximating the original vectors in a manner that the cosine of the angles is nearly preserved:

$$\cos(\theta) = \frac{\Delta I(p) \cdot \Delta I(q)}{\|\Delta I(p)\| \|\Delta I(q)\|} \approx \frac{\Delta T(p) \cdot \Delta T(q)}{\|\Delta T(p)\| \|\Delta T(q)\|}. \quad (13)$$

Since T approximates Φ , and Φ is conformal, T aims to preserve the cosines of angles between points in D , thus achieving an approximation of conformal mapping in the discrete pixel space. MPD utilizes this property of T to maintain the visual integrity of the original image I in the transformed image I' , where $I'(p) = T(p)$ for all $p \in D$.

7 Perspectively distorted views synthesis through MPD

MPD distorts the original image to synthesize perspective distortion based on steps shown in Algorithm 1 in main paper, resulting in a distorted image with a black background (*variant 1*, Fig. 15). Optionally, to remove the black background, replication padding is used to obtain integrated padding variant (*variant 2*, Fig. 15) of transformed image in step 9 of Algorithm 1.



Fig. 15: MPD transformed images - (black background) variant 1 and (integrated padding background) variant 2

In connection to Fig. 1, we additionally show the integrated padding background variants of different orientations views synthesizing PD in Fig. 16.



Fig. 16: MPD IB (integrated padding background) synthesizes perspective distortion with different orientations corresponding to the parameters. It replaces the black background with boundary foreground pixel values of the transformed image.

Fig. 17 Demonstrate controlled scaling of distortion in MPD-transformed image for different views in detail. It showcases the capabilities of MPD to mimic real-world distortions synthetically, which are otherwise challenging to capture and train the models.

8 Exploration of parameters a , b , and d

Fig. 18 shows the effect of parameters a , b , and d , representing subsets of affine transformations.

9 MPD in self-supervised learning methods

We present a comprehensive investigation into the application of MPD (Geometry-Exploring Encoded Transformation) within self-supervised learning (SSL) approaches, with a primary focus on its integration with the contrastive learning method SimCLR [7], shown in Fig. 19. Our core objective is to assess MPD’s effectiveness in enhancing SSL’s representation learning capacities by adapting to perspective variations. Section 3.3 in the main paper explains the detailed results of MPD incorporation with simCLR contrastive learning to mitigate perspective distortion. Further Table linear evaluation results are included in Section 12. Furthermore, we extend our exploration to evaluate MPD’s compatibility and performance in conjunction with another SSL method, DINO [6], which is grounded in knowledge distillation-based self-supervised approaches. Results on MPD incorporation in DINO (Fig. 20) are presented in Table 15 and Fig. 22 in Section 12. **Probability of 0.8:** Incorporating MPD with a probability of 0.8 within the SimCLR contrastive learning shows a well-balanced strategy. This allocation ensures that each of the four perspective views—left, right, top, and bottom—receives an equitable opportunity approximately 20% of the time, fostering diversity and divergence within the training data. This design navigates the intricate terrain of the diversity-consistency trade-off with precision, enabling our model to gain insights from multiple angles while maintaining the stability and consistency garnered

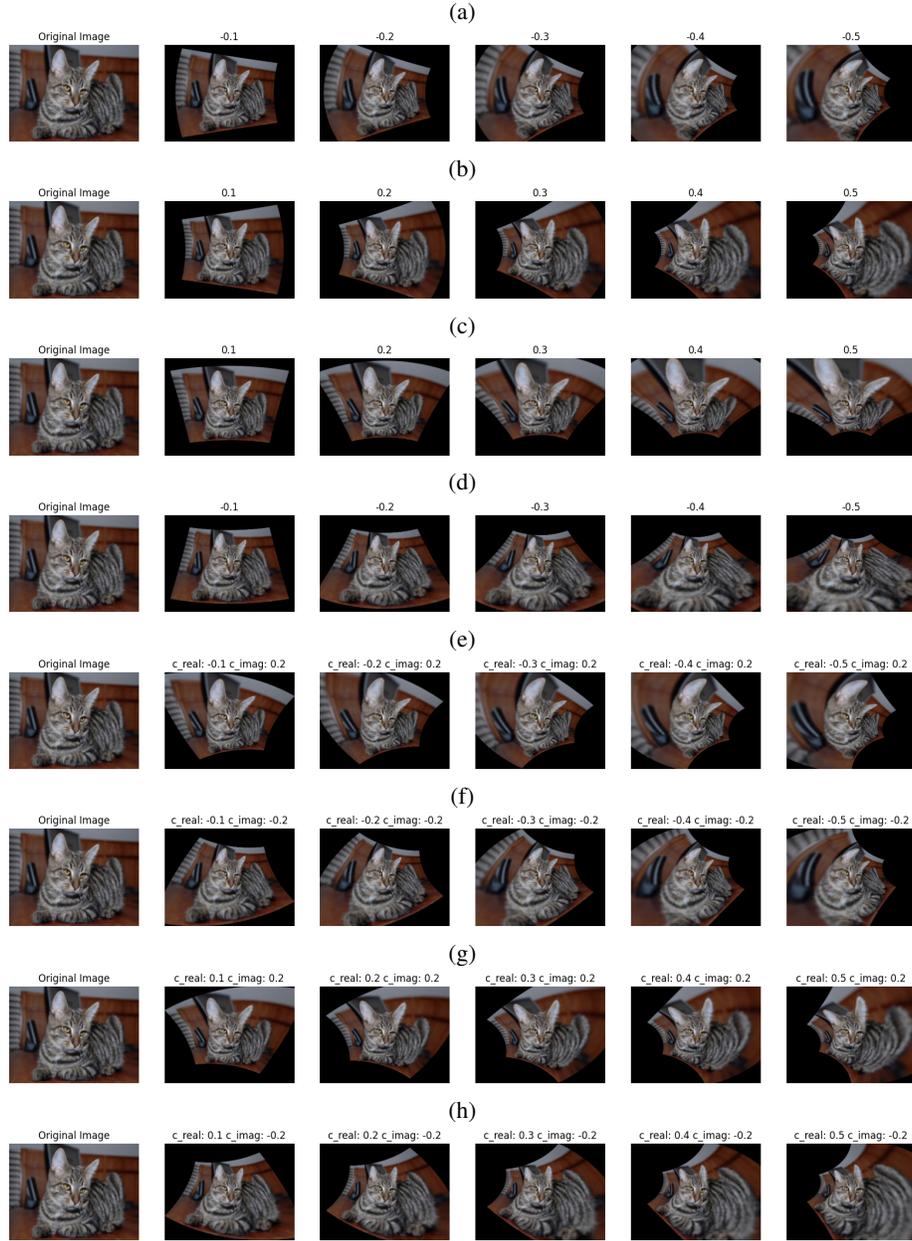


Fig. 17: Demonstrate controlled scaling of distortion in MPD-transformed images for different views. a) distorted left view by varying c_{real} from -0.1 to -0.5 b) distorted right view by varying c_{real} from $+0.1$ to $+0.5$ c) distorted top view by varying c_{imag} from $+0.1$ to $+0.5$ d) distorted bottom view by varying c_{imag} from -0.1 to -0.5 e) distorted left-top view by varying c_{real} from -0.1 to -0.5 with fixed c_{imag} to $+0.2$ f) distorted left-bottom view by varying c_{real} from -0.1 to -0.5 with fixed c_{imag} to -0.2 g) distorted right-top view by varying c_{real} from $+0.1$ to $+0.5$ with fixed c_{imag} to $+0.2$ h) distorted right-bottom view by varying c_{real} from $+0.1$ to $+0.5$ with fixed c_{imag} to -0.2 . Current work experiments only use left, right, top, and bottom view synthesis.

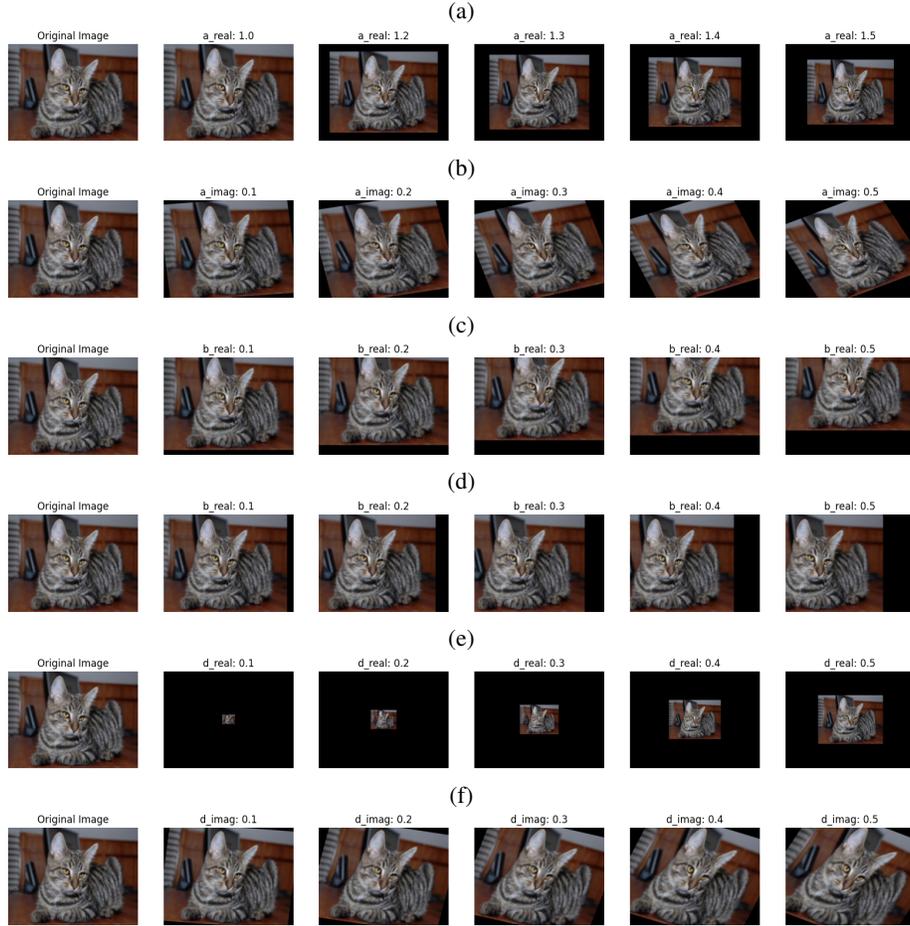


Fig. 18: Visualizing the effect of parameters a , b , and d , representing subsets of affine transformations. a) imitating resizing by controlling parameter (increasing value decreases size) a_{real} 1.0 to 1.5 b) rotating (anti-clockwise) by controlling parameter a_{imag} 0.1 to 0.5 c) imitating vertical translation by controlling parameter b_{real} 0.1 to 0.5 d) imitating horizontal translation by controlling parameter b_{imag} 0.1 to 0.5 e) imitating resizing by controlling parameter (increasing value increase size) d_{real} 0.1 to 0.5 f) imitating rotation by controlling parameter (clockwise) d_{imag} 0.1 to 0.5

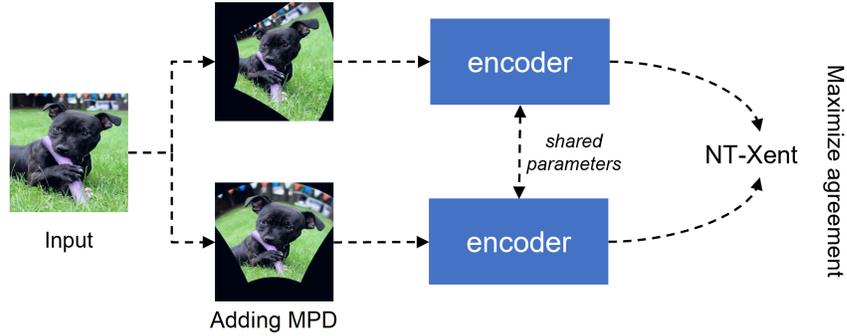


Fig. 19: ssl:MPD - MPD integrated in SimCLR [7] contrastive learning self-supervised method. MPD introduces perspective distortion by transforming the views of the input image. Standard SimCLR and ssl:DINO-MPD both are pretrained for 100 epochs with a batch size of 512 and linear evaluation for 100 epochs with a batch size of 256. ResNet50 is backbone.

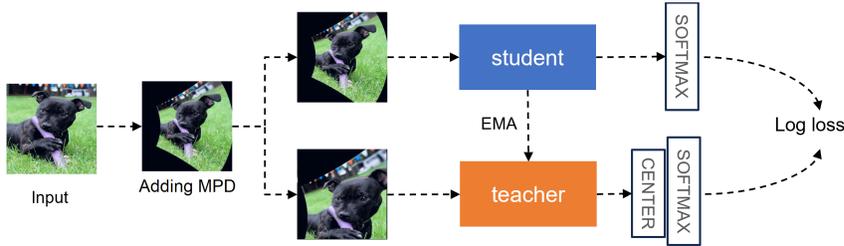


Fig. 20: ssl:DINO-MPD - MPD integrated in DINO [6] knowledge distillation based self-supervised method. MPD introduces perspective distortion by transforming the input image before the creation of views. Standard DINO and ssl:DINO-MPD both are pretrained for 100 epochs with a batch size of 512 and linear evaluation for 100 epochs with a batch size of 256. ViT-small is backbone.

from the original, unaltered views. The empirical validation of our approach underscores its efficacy, as it not only enhances the model’s ability to recognize objects and scenes from varying perspectives but also fortifies its robustness against real-world perspective distortions. The DINO method applies MPD to the original input image before creating local and global views (as defined in the DINO method), so the MPD probability is kept at 0.4. All the results described with standard deviation are repeated three times.

10 Continuation (Sec. 3.1): Existing models’ robustness evaluation

Extended evaluations on the ImageNet-PD benchmark dataset, detailed in Tables 7 and 8, reveal how different architectures respond to perspective distortions, extending the main paper’s insights (Fig. 4). Models like ResNet50 and ViT-L-16 show varied robustness to distortions such as Left-view (PD-L and PD-LI), with ResNet50 maintain-

ing 63.37% Top-1 accuracy for PD-L compared to 61.15% for PD-B, and VIT-L-16 at 74.34% for PD-LI versus 71.73% for PD-BI. With their self-attention mechanisms, Vision Transformers demonstrate a smaller performance drop than traditional CNNs under perspective distortions, indicating their potential to maintain representational robustness. As discussed in our main paper, this analysis underscores the need for adaptive, perspective-aware models like MPD in real-world applications.

Table 7: Detailed results of ImageNet trained models with standard CNN and transformer architectures; evaluated on original validation set and ImageNet-PD subsets with black background. Results in continuation of Fig. 4 (first plot).

Model	Original Validation Set		Perspectively Distorted Top-view (PD-T)		Perspectively Distorted Bottom-view (PD-B)		Perspectively Distorted Left-view (PD-L)		Perspectively Distorted Right-view (PD-R)		PD Average	
	Top1	Top5	Top1	Top5	Top1	Top5	Top1	Top5	Top1	Top5	Top1	Top5
AlexNet	56.52	79.07	36.74	60.02	37.08	61.21	37.87	61.87	37.33	61.06	37.26	61.04
EfficientNet-b2	80.61	95.31	67.14	86.49	64.97	84.84	67.60	87.11	67.15	86.77	66.71	86.30
GoogleNet	69.78	89.53	54.63	76.97	55.29	78.45	55.39	78.24	54.68	77.72	55.00	77.85
Inception-v3	77.29	93.45	55.09	76.76	52.79	74.83	54.45	76.19	54.86	76.58	54.30	76.09
MobileNet-v2	71.88	90.29	57.60	79.78	56.26	78.73	57.40	79.91	56.87	79.54	57.03	79.49
ResNet101	77.37	93.55	65.86	85.58	65.17	85.29	65.62	85.46	65.40	85.49	65.51	85.46
ResNet50	76.13	92.86	63.37	83.61	61.15	81.86	65.20	85.13	65.84	85.64	63.89	84.06
ResNet50 (SSL)	76.36	92.99	63.76	84.12	61.71	82.91	65.20	85.13	65.96	86.11	64.16	84.57
ResNet34	73.31	91.42	59.23	80.84	58.65	80.60	59.72	81.25	58.79	80.76	59.10	80.86
ResNet18	69.76	89.08	56.00	78.52	55.07	77.90	55.68	78.73	55.27	78.01	55.50	78.29
VGG19	72.38	90.88	58.50	80.36	58.09	80.36	58.00	80.13	57.64	79.67	58.06	80.13
VGG16	71.59	90.38	57.94	80.54	57.32	79.89	57.47	79.91	57.75	80.15	57.62	80.12
VIT-B-16	81.07	95.32	73.94	90.87	73.63	90.73	74.50	91.31	74.30	91.06	74.09	90.99
VIT-L-16	79.66	94.64	74.05	91.06	74.30	90.97	74.34	91.08	73.93	90.91	74.15	91.00

Table 8: Detailed results of ImageNet trained models with standard CNN and transformer architectures; evaluated on original validation set and ImageNet-PD subsets with integrated padding background. Results in continuation of Fig. 4 (second plot).

Model	Original Validation Set		Perspectively Distorted Top-view (PD-TI)		Perspectively Distorted Bottom-view (PD-BI)		Perspectively Distorted Left-view (PD-LI)		Perspectively Distorted Right-view (PD-RI)		PD Average	
	Top1	Top5	Top1	Top5	Top1	Top5	Top1	Top5	Top1	Top5	Top1	Top5
AlexNet	56.52	79.07	40.73	64.47	41.09	65.36	42.53	66.32	41.13	65.08	41.37	65.31
EfficientNet-b2	80.61	95.31	67.83	87.28	66.29	85.95	68.16	87.43	68.12	87.48	67.60	87.04
GoogleNet	69.78	89.53	54.68	77.22	56.84	79.51	57.24	79.68	56.01	78.82	56.19	78.81
Inception-v3	77.29	93.45	56.77	78.36	54.32	76.23	55.94	77.56	56.81	78.55	55.96	77.68
MobileNet-v2	71.88	90.29	57.34	79.68	57.84	80.20	58.18	80.42	57.26	79.96	57.66	80.07
ResNet101	77.37	93.55	64.30	84.47	65.57	85.64	66.55	86.14	65.69	85.70	65.53	85.49
ResNet50	76.13	92.86	62.26	83.07	63.39	83.91	63.81	84.17	63.57	84.16	63.26	83.83
ResNet50 (SSL)	76.24	92.90	61.73	82.50	61.91	83.07	65.58	85.74	65.79	86.11	63.75	84.36
ResNet34	73.31	91.42	58.29	80.11	59.61	81.38	60.65	82.22	59.20	81.06	59.44	81.19
ResNet18	69.76	89.08	55.85	78.58	56.27	79.13	56.87	79.83	56.05	78.57	56.26	79.03
VGG19	72.38	90.88	57.68	80.06	59.28	81.43	58.05	80.41	57.58	79.73	58.15	80.41
VGG16	71.59	90.38	56.26	78.97	58.58	80.94	57.96	80.37	57.41	80.07	57.55	80.09
VIT-B-16	81.07	95.32	70.49	88.58	71.89	89.49	72.45	89.93	71.51	89.27	71.59	89.32
VIT-L-16	79.66	94.64	70.71	88.73	71.73	89.41	71.87	89.49	71.25	89.16	71.39	89.20

11 Continuation (section 3.2): MPD’s effects on supervised learning

In the extended analysis presented in Tables 9 and 10, we observe a shift in the performance of MPD and MPD IB when applied at a probability of 1.0, approximately following the trends reported in the main paper up to a probability of 0.8 (Fig. 9). For

MPD at $P=1.0$, a marginal decrease in Top-1 accuracy on the original ImageNet validation set ($74.10\% \pm 0.03$) compared to $P=0.8$ ($76.34\% \pm 0.02$), the robustness against perspective distortions on ImageNet-PD subsets is maintained, with accuracies between $71.57\% \pm 0.04$ to $72.05\% \pm 0.04$. This highlights a nuanced balance between handling distortions and achieving peak performance on standard datasets. Similarly, MPD IB at $P=1.0$ follows a comparable pattern. The Top-1 accuracy on the standard ImageNet set slightly reduces to $74.62\% \pm 0.03$ from $75.63\% \pm 0.05$ at $P=0.8$. However, its performance on the ImageNet-PD subsets, though marginally lower than MPD, remains consistent with Top-1 accuracies in the range of $68.23\% \pm 0.04$ to $70.30\% \pm 0.03$. These results emphasize the trade-offs involved in fully applying MPD and MPD IB and the importance of finding a balanced approach to optimize performance across both standard data and perspective distortion present in data.

ImageNet-E: In our extended analysis of MPD and MPD IB within a supervised approach (Table 11 and Fig. 6), we observe notable performance improvements in size change scenarios of ImageNet-E, indicative of enhanced robustness against perspective distortion. These models, specifically in size-related subsets (full, 0.10, 0.08, 0.05), show a reduced drop in Top-1 accuracy and higher absolute accuracy than standard ResNet50. However, in background change subsets, MPD and MPD IB’s performance occasionally falls behind ResNet50, reflecting their focus on perspective distortions over background variations. This further aligns with the findings on MPD’s capabilities in perspective-centric scenarios.

ImageNet-X Detailed analysis in Table 12 extends the findings from the main paper (Fig. 8), showcasing the superior performance of supervised MPD models on ImageNet-X, especially in managing perspective-related factors. These models demonstrate a marked improvement in handling size variations (‘larger’ and ‘smaller’) and complex scenarios involving ‘object blocking’ and ‘person blocking’ compared to the standard ResNet50. The error ratio and absolute accuracy across these factors highlight the robustness of MPD in dealing with perspective challenges, reflecting the conclusion about MPD’s enhanced learning capabilities in diverse visual conditions.

12 Continuation (section 3.3): MPD’s effects on self-supervised learning

SimCLR [7] self-supervised MPD models demonstrate notable effectiveness on ImageNet-E and ImageNet-X, detailed in Tables 11 and 12. On ImageNet-E, self-supervised MPD variants show a marked reduction in error ratios in size-related subsets, with error ratios like 1.44 for ‘larger’ and 1.46 for ‘smaller’, indicating a strong grasp of perspective challenges. In ImageNet-X, these models excel in handling ‘object blocking’ (error ratio 1.61) and ‘person blocking’ (error ratio 1.97) factors, achieving higher absolute accuracy compared to standard ResNet50, as illustrated in Fig. 8. The detailed analysis in Fig. 21 further emphasizes the models’ enhanced adaptability in self-supervised supervised learning, underlining their robustness in understanding complex visual scenarios.

Linear evaluation results of SimCLR [7] self-supervised MPD (ssl:MPD) and MPD IB (ssl:MPD IB) models, presented in Tables 13 and 14, align remarkably with the trends observed in their fine-tuned counterparts (Tables 7 and 8). This complementary relation-

Table 9: Extended comparative analysis of MPD trained on supervised (*supervised:MPD model*) and self-supervised (*ssl:MPD model*) approaches and evaluated on original ImageNet and **ImageNet-PD subsets** (black background). The probability P of applying MPD in model fine-tuning is in the first column. Added results for probability values 0.0 and 1.0 in continuation to Table 1.

P	Original Validation Set		Perspectively Distorted Top-view (PD-T)		Perspectively Distorted Bottom-view (PD-B)		Perspectively Distorted Left-view (PD-L)		Perspectively Distorted Right-view (PD-R)	
	Top1	Top 5	Top1	Top 5	Top1	Top 5	Top1	Top 5	Top1	Top 5
Supervised training from scratch										
0.0	76.13±0.04	92.86±0.01	63.37±0.06	83.61±0.02	61.15±0.04	81.86±0.01	65.20±0.03	85.13±0.03	65.84±0.06	85.64±0.02
0.2	76.05±0.03	92.99±0.02	72.48±0.02	91.02±0.01	72.12±0.02	91.08±0.02	72.57±0.03	91.13±0.03	72.80±0.05	91.35±0.02
0.4	76.17±0.04	93.03±0.02	73.13±0.03	91.33±0.01	72.94±0.02	91.42±0.01	73.19±0.06	91.48±0.01	73.38±0.01	91.69±0.01
0.6	76.19±0.05	93.14±0.03	73.23±0.02	91.46±0.03	73.01±0.05	91.34±0.02	73.47±0.02	91.66±0.01	73.54±0.06	91.61±0.02
0.8	76.34±0.02	93.03±0.02	73.00±0.02	90.69±0.02	72.31±0.03	90.81±0.03	73.50±0.04	91.33±0.03	72.91±0.02	91.29±0.02
1.0	74.10±0.03	91.50±0.03	71.57±0.04	90.31±0.03	71.65±0.03	90.44±0.02	71.88±0.04	90.62±0.03	72.05±0.04	90.60±0.03
Self-supervised pre-training on contrastive learning integrating MPD (probability for pre-training=0.8)										
0.2	76.37±0.05	93.60±0.02	72.34±0.02	90.81±0.02	72.24±0.03	90.95±0.02	72.48±0.04	91.08±0.01	72.81±0.03	91.16±0.01
0.4	76.77±0.02	93.40±0.01	73.23±0.03	91.39±0.01	73.06±0.04	91.54±0.02	73.55±0.03	91.53±0.03	73.54±0.05	91.65±0.03
0.6	76.14±0.04	93.58±0.02	73.39±0.02	91.50±0.02	73.29±0.03	91.57±0.01	73.57±0.03	91.66±0.02	73.73±0.04	91.58±0.02
0.8	76.29±0.03	92.78±0.02	73.61±0.04	91.49±0.03	73.27±0.05	91.48±0.03	73.60±0.06	91.69±0.01	73.81±0.02	91.71±0.02
1.0	74.13±0.04	91.89±0.03	71.70±0.03	90.42±0.03	71.79±0.04	90.65±0.02	71.96±0.04	90.70±0.02	72.20±0.03	90.88±0.02

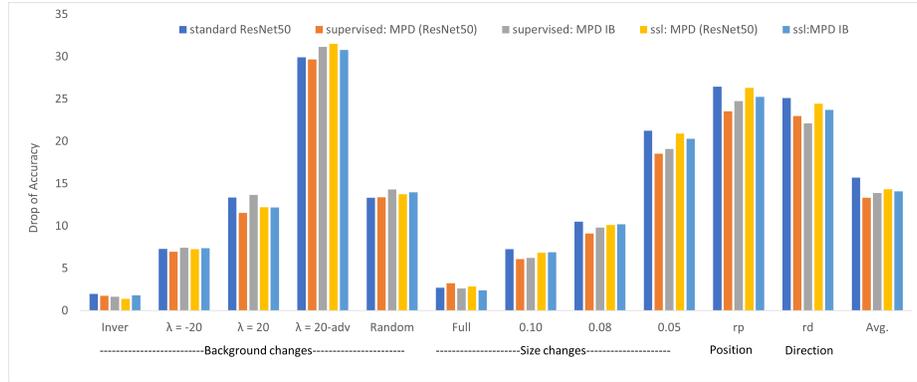


Fig. 21: Detailed performance comparison on ImageNet-E [28]: Detailed results on Drop of Top1 accuracy under background changes, size changes, random position (rp), random direction (rd), and average over 11 subsets. Lower is better. *Average* reports the mean of all subsets. MPD variants outperforms standard ResNet50 on subsets which are affected with perspective distortion, e.g., size, position, and direction. Corresponding details on absolute accuracy measure are in Table 11. Detailed analysis in continuation of Fig. 6; Fig. 6 is given in main paper.

ship underscores the robustness of MPD models in both linear and fine-tuned scenarios. Specifically, ssl:MPD and ssl:MPD IB demonstrate impressive handling of perspective distortion in the ImageNet-PD subsets during linear evaluation, with ssl:MPD achieving an average Top-1 accuracy of $50.33±0.04$, and ssl:MPD IB even more impressive at $51.62±0.03$. These figures notably surpass the standard self-supervised model’s performance, which scores $35.08±0.04$ and $32.53±0.04$, respectively. This substantial improvement in handling perspective distortions is consistent with the fine-tuned results, where both ssl:MPD and ssl:MPD IB models maintained high accuracy levels, show-

Table 10: Extended comparative analysis of MPD IB trained on supervised (*supervised:MPD IB model*) and self-supervised (*ssl:MPD IB model*) approaches and evaluated on original ImageNet and **ImageNet-PD subsets** (black background). The probability P of applying MPD BI in model fine-tuning is in the first column.

P	Original Validation Set		Perspectively Distorted Top-view (PD-TI)		Perspectively Distorted Bottom-view (PD-BI)		Perspectively Distorted Left-view (PD-LI)		Perspectively Distorted Right-view (PD-RI)	
	Top1	Top 5	Top1	Top 5	Top1	Top 5	Top1	Top 5	Top1	Top 5
Supervised training from scratch										
0.0	76.13±0.04	92.86±0.01	61.59±0.04	82.39±0.03	61.80±0.05	82.40±0.02	65.30±0.05	85.09±0.02	66.11±0.05	85.92±0.02
0.2	75.78±0.04	92.76±0.02	71.68±0.01	90.48±0.03	71.83±0.04	90.48±0.02	72.09±0.04	90.73±0.01	72.26±0.05	90.90±0.01
0.4	75.88±0.03	92.92±0.03	72.37±0.04	90.35±0.02	72.31±0.04	90.78±0.03	72.43±0.05	91.12±0.03	72.77±0.05	91.09±0.03
0.6	76.14±0.06	93.05±0.03	72.99±0.03	90.78±0.03	72.59±0.04	91.24±0.03	73.11±0.03	91.41±0.02	73.42±0.01	91.52±0.01
0.8	75.63±0.05	92.73±0.02	73.11±0.06	91.23±0.02	73.02±0.02	91.43±0.02	73.33±0.02	91.62±0.02	73.51±0.06	91.50±0.01
1.0	74.62±0.03	91.88±0.02	68.23±0.04	87.80±0.02	69.02±0.02	88.48±0.03	69.83±0.03	89.19±0.03	70.30±0.03	89.39±0.02
Self-supervised pre-taining on contrastive learning integrating MPD (probability for pre-training=0.8)										
0.2	75.43±0.02	92.59±0.01	72.38±0.06	90.71±0.02	72.26±0.04	90.83±0.03	72.40±0.05	90.88±0.01	72.61±0.04	90.91±0.02
0.4	75.58±0.03	92.67±0.01	72.87±0.01	91.28±0.01	73.17±0.04	91.39±0.02	73.34±0.03	91.58±0.03	73.41±0.02	91.49±0.03
0.6	75.41±0.04	92.64±0.03	73.43±0.06	91.41±0.01	73.01±0.05	91.50±0.02	73.37±0.02	91.64±0.01	73.61±0.05	91.60±0.02
0.8	75.23±0.01	92.46±0.01	73.54±0.03	91.59±0.02	73.41±0.02	91.57±0.01	73.60±0.03	91.73±0.02	73.79±0.04	91.84±0.02
1.0	74.98±0.02	91.35±0.02	68.44±0.03	88.10±0.03	69.46±0.01	89.05±0.02	70.08±0.03	89.88±0.03	70.88±0.02	90.24±0.02

Table 11: Detailed results on ImageNet-E: (a) Drop of Top1 accuracy (Lower is better) and (b) Absolute Accuracy (Higher is better) under background changes, size changes, random position (rp), random direction (rd), and average over 11 subsets.

(a) Drop of Accuracy													
Models	Original	Background changes					Size changes			Position	Direction	Avg.	
		Inver	$\lambda = -20$	$\lambda = 20$	$\lambda = 20$ -adv	Random	Full	0.10	0.08	0.05	rp		rd
standard ResNet50	92.69	1.97	7.30	13.35	29.92	13.34	2.71	7.25	10.51	21.26	26.46	25.12	15.72
supervised: MPD (ResNet50)	92.51	1.75	6.94	11.53	29.66	13.40	3.22	6.07	9.11	18.54	23.55	22.98	13.34
supervised: MPD IB	92.39	1.63	7.42	13.67	31.16	14.32	2.62	6.23	9.82	19.09	24.75	22.10	13.89
ssl: MPD (ResNet50)	92.51	1.38	7.26	12.22	31.53	13.76	2.85	6.85	10.12	20.91	26.31	24.43	14.33
ssl:MPD IB	92.30	1.82	7.38	12.18	30.79	13.97	2.40	6.90	10.21	20.29	25.25	23.71	14.08
(b) Absolute Accuracy													
Models	Original	Background changes					Size changes			Position	Direction	Avg.	
		Inver	$\lambda = -20$	$\lambda = 20$	$\lambda = 20$ -adv	Random	Full	0.10	0.08	0.05	rp		rd
standard ResNet50	92.69	90.72	85.39	79.34	62.77	79.35	89.98	85.44	82.18	71.43	66.23	67.57	78.22
supervised: MPD (ResNet50)	92.51	90.76	85.57	80.97	62.84	79.11	89.29	86.44	83.40	73.97	68.96	69.53	79.17
supervised: MPD IB	92.39	90.76	84.97	78.72	61.24	78.08	89.77	86.16	82.57	73.30	67.65	70.29	78.50
ssl: MPD (ResNet50)	92.51	91.13	85.25	80.28	60.98	78.75	89.66	85.66	82.39	71.60	66.20	68.08	78.18
ssl:MPD IB	92.30	90.49	84.93	80.12	61.51	78.33	89.90	85.40	82.09	72.01	67.05	68.59	78.22

Table 12: Detailed results on ImageNet-X [23]: (a) Error ratio (close to 1.0 is best) and (b) Absolute accuracy (higher is better) of 16 factors. *Average* reports the mean accuracy of all factors. Detailed results also includes *supervised:MPD IB* and *ssl:MPD IB* models besides *supervised:MPD* and *ssl:MPD* models. All the models achieves improved error ratio with higher accuracy consistently across factors. Detailed results in continuation of Fig. 8; Fig. 8 is given in main paper.

(a) Error Ratio					
Factor	standard ResNet50	supervised: MPD	ssl: MPD	supervised: MPD IB	ssl: MPD IB
pose	0.70	0.72	0.74	0.73	0.72
pattern	0.77	0.89	0.78	0.87	0.87
partial_view	1.13	0.89	0.86	1.07	0.90
background	1.14	1.08	1.09	1.09	1.07
brighter	1.17	1.09	1.09	1.19	1.13
color	1.19	1.21	1.18	1.23	1.20
larger	1.32	1.30	1.22	1.24	1.35
smaller	1.50	1.44	1.50	1.46	1.50
darker	1.72	1.52	1.51	1.46	1.53
shape	1.75	1.63	1.61	1.53	1.59
object_blocking	1.92	1.66	1.61	1.75	1.62
style	1.93	1.73	1.63	1.80	1.71
multiple_objects	1.97	1.76	1.72	1.84	1.86
subcategory	2.10	1.92	1.90	1.88	1.98
person_blocking	2.17	1.95	1.97	1.94	2.03
texture	2.39	2.20	2.18	2.21	2.16
Avg	1.55	1.44	1.41	1.45	1.45
(b) Absolute Accuracy					
Factor	standard ResNet50	supervised: MPD	ssl: MPD	supervised: MPD IB	ssl: MPD IB
pose	82.40	81.47	80.55	81.12	81.00
pattern	80.48	77.28	79.55	77.41	77.27
partial_view	71.32	77.27	77.49	72.41	76.24
background	71.13	72.37	71.53	71.87	71.85
brighter	70.45	72.10	71.32	69.30	70.38
color	69.85	69.12	68.96	68.18	68.44
larger	66.67	66.67	68.00	68.00	64.67
smaller	62.07	63.11	60.62	62.30	60.66
darker	56.56	61.17	60.47	62.28	59.93
shape	55.67	58.32	57.88	60.47	58.32
object_blocking	51.28	57.50	57.69	54.64	57.50
style	51.16	55.81	57.38	53.33	55.13
multiple_objects	50.00	55.00	55.00	52.50	51.16
subcategory	46.84	50.94	50.09	51.28	48.03
person_blocking	45.00	50.00	48.33	49.91	46.67
texture	39.36	43.62	42.91	42.91	43.26
Avg	60.64	63.23	62.99	62.37	61.91

ing their efficacy in learning robust representations. Furthermore, the linear evaluation results reveal that these improvements are not at the expense of general performance. Both *ssl:MPD* and *ssl:MPD IB* sustain competitive accuracies on the standard ImageNet validation set, mirroring the fine-tuning outcomes where they demonstrated broad applicability and versatility. These findings from the linear evaluation and fine-tuning phases collectively reinforce the conclusion in our main paper about the enhanced capabilities of MPD models in self-supervised learning environments. They affirm the models’ ability to learn perspective-aware representations that are not only applicable to specialized scenarios involving perspective distortions but also effective across a wide range of visual recognition tasks. Linear evaluation results for *ssl:MPD* and *ssl:MPD IB* strongly reinforce the claim made in our main paper about their enhanced ability in self-supervised learning contexts. These models show not just incremental but sig-

nificant improvements in handling complex visual distortions, establishing MPD as a powerful tool for self-supervised learning in real-world applications where perspective variations are prevalent.

Table 13: Linear evaluation performance comparison for self-supervised models, standard SimCLR [7] and *ssl:MPD* model on original ImageNet validation set and ImageNet-PD subsets (*black background*). Self-supervised pretraining was performed on ImageNet dataset with batch size of 512 on linear learning rate for 100 epochs, refer SimCLR [7]. Top-1 accuracy reported. The probability of MPD is set to 0.8 for self-supervised pre-training and linear evaluation. ResNet50 is common backbone.

Model	Original Validation Set	Perspectively Distorted	Perspectively Distorted	Perspectively Distorted	Perspectively Distorted
		Top-view (PD-T)	Bottom-view (PD-B)	Left-view (PD-L)	Right-view (PD-R)
standard SimCLR	60.14	34.22	32.87	36.33	36.89
<i>ssl:MPD</i>	60.02±0.03	50.63±0.03	49.65±0.03	50.41±0.03	50.64±0.05

Table 14: Linear evaluation performance comparison for self-supervised models, standard SimCLR [7] and *ssl:MPD IB* model on original ImageNet validation set and ImageNet-PD subsets (*with integrated padding background*). Self-supervised pretraining was performed on batch size of 512 on linear learning rate for 100 epochs, refer SimCLR [7]. Top-1 accuracy reported. The probability of MPD IB is set to 0.8 for self-supervised pre-training and linear evaluation. ResNet50 is common backbone.

Model	Original Validation Set	Perspectively Distorted	Perspectively Distorted	Perspectively Distorted	Perspectively Distorted
		Top-view (PD-TI)	Bottom-view (PD-BI)	Left-view (PD-LI)	Right-view (PD-RI)
standard SimCLR	60.14	30.41	30.18	34.81	34.72
<i>ssl:MPD IB</i>	60.05±0.04	51.30±0.03	51.45±0.04	51.81±0.03	51.93±0.03

MPD with DINO: Incorporating MPD into the knowledge distillation-based self-supervised method DINO [6] demonstrates MPD’s adaptability and effectiveness across different SSL approaches. The linear evaluation results for *ssl:DINO-MPD* and *ssl:DINO-MPD IB*, as detailed results presented in Tables 15 and 16, demonstrate the significant impact of integrating MPD and MPD IB into the DINO method. Both *ssl:DINO-MPD* and *ssl:DINO-MPD IB* models substantially improves robustness against perspective distortion when evaluated on the ImageNet-PD subsets. Specifically, *ssl:DINO-MPD* achieves an average Top-1 accuracy of 60.70 ± 0.04 , while *ssl:DINO-MPD IB* further enhances this performance, achieving an even higher average Top-1 accuracy of 61.30 ± 0.02 . The integration of MPD and MPD IB into DINO demonstrates their adaptability and effectiveness in improving performance across self-supervised approaches, especially in scenarios involving perspective distortion.

Furthermore, the activation maps from *ssl:DINO-MPD*, as visualized in Fig. 22, qualitatively demonstrate its superior understanding of perspective distortion in input images. These maps reveal a more focused and relevant visual attention, indicating a deeper comprehension of the image structure and content, especially in perspective distorted scenarios. The success of *ssl:DINO-MPD* in both quantitative performance and qualitative visual attention underscores the versatility of MPD. It shows that

Table 15: Linear evaluation performance comparison for knowledge distillation based self-supervised method DINO [6]. Standard DINO is original method pretrained for 100 epochs with batch size of 512 and *ssl: DINO-MPD* model is MPD integrated DINO, following same pre-training hyperparameters. Post-pretraining, both the models are linearly trained for 100 epochs with batch size of 256. Top-1 accuracy reported. Evaluated on original ImageNet validation set and ImageNet-PD subsets (*black background*). The probability of MPD is set to 0.4 for self-supervised pre-training and 0.8 for linear evaluation. ViT-small transformer [13] is common backbone.

Model	Original Validation Set	Perspectively Distorted	Perspectively Distorted	Perspectively Distorted	Perspectively Distorted
		Top-view (PD-T)	Bottom-view (PD-B)	Left-view (PD-L)	Right-view (PD-R)
standard DINO	74.00	46.31	46.05	46.15	45.98
ssl:DINO-MPD	72.36±0.01	60.72±0.02	60.86±0.02	60.63±0.02	60.58±0.02

Table 16: Linear evaluation performance comparison for DINO-MPD IB: Standard DINO [6] is original method pretrained for 100 epochs with batch size of 512 and *ssl: DINO-MPD IB model* is MPD IB integrated DINO, following same pretraining hyperparameters. Post-pretraining, both the models are linearly trained for 100 epochs with batch size of 256. Top-1 accuracy reported. Evaluated on original ImageNet validation set and ImageNet-PD subsets (*with integrated padding background*). The probability of MPD IB is set to 0.4 for self-supervised pre-training and 0.8 for linear evaluation. ViT-small transformer [13] is common backbone.

Model	Original Validation Set	Perspectively Distorted	Perspectively Distorted	Perspectively Distorted	Perspectively Distorted
		Top-view (PD-T)	Bottom-view (PD-B)	Left-view (PD-L)	Right-view (PD-R)
standard DINO	74.00	45.22	44.80	44.05	45.66
ssl:DINO-MPD IB	72.05±0.01	61.04±0.01	61.50±0.02	60.98±0.03	61.66±0.02

MPD can be effectively combined with various SSL methods, preserving their inherent strengths and augmenting them to handle better the challenges posed by perspective distortions. This enhances the claims made in our main paper about MPD’s capabilities, providing a strong case for its widespread applicability in diverse self-supervised learning approaches.

13 MPD vs. other augmentations

Table 17 demonstrates the effectiveness of MPD against other popular augmentation methods on the ImageNet validation set, ImageNet-PD subsets, and ImageNet-E [28]. Augmentation methods such as Mixup [58], Cutout [12], AugMix [17], and Pixmix [18] showed slight improvements on the ImageNet-PD benchmarks. For instance, Mixup improved the Top-1 accuracy to 65.46% on PD-T and 68.43% on PD-R, but these methods generally degraded performance on ImageNet-E, with Mixup dropping to 82.38%. This suggests that these augmentations do not adequately capture the characteristics of perspective distortions. In contrast, MPD-trained models, both in supervised and self-supervised settings, consistently outperformed all other methods on the perspective distortion benchmarks and ImageNet-E. The supervised MPD model achieved impressive Top-1 accuracies of 73.00%, 72.31%, 73.50%, and 72.91% on PD-T, PD-B, PD-L, and PD-R, respectively, and 85.92% on ImageNet-E. The self-supervised MPD model further enhanced performance, achieving Top-1 accuracies of 73.23%, 73.06%, 73.55%, and 73.54% on PD-T, PD-B, PD-L, and PD-R, respectively, and 86.66% on

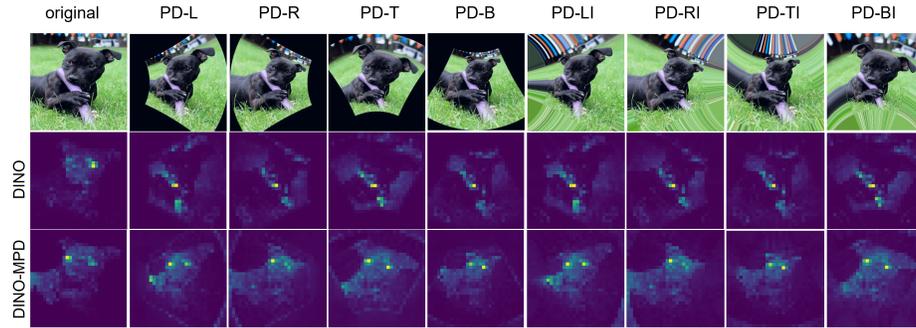


Fig. 22: Activation maps of visual attention: Comparison of attention maps between DINO and MPD integrated DINO (*ssl:DINO-MPD* model). First attention head is used to visualize the attention. Example taken from ImageNet-PD set with original image. Both the models are pretrained for 100 epochs using ViT small architecture. DINO-MPD demonstrates on learning perspective distortion in input image.

ImageNet-E. These results indicate that MPD effectively addresses perspective distortions, maintaining high accuracy across both the original and distorted datasets.

14 Continuation (section 3.4): MPD’s generalizability

This section extends the results for MPD’s adaptability in Fisheye, Crowd counting, and Person re-identification.

14.1 Crowd counting

ResNet50 encoder is employed in P2P-Net [45], the integration of MPD and MPD-AutoCrowd continue to demonstrate remarkable improvements, shown in Table 18. For instance, the standard P2P-Net with ResNet50 records an MAE of 299.5 and an MSE of 382.48 on the UCF-CC50 dataset. In contrast, the incorporation of MPD significantly enhances performance, yielding an MAE of 95.10 and an MSE of 132.47. Moreover, MPD-AutoCrowd further improves these metrics, achieving an MAE of 105.4 and an MSE of 124.61. These results underscore the capability of MPD to effectively deal with the complexities of crowd scenes, even when coupled with a different encoder. This ablation study complements the findings presented in Table 2 from the main paper, where MPD-CC and MPD-AutoCrowd outperform other crowd-counting methods across various benchmarks. The consistent improvement across different datasets and encoder configurations, especially with the challenging UCF-CC50 dataset, highlights MPD’s robustness and adaptability. MPD-AutoCrowd’s notable performance in handling varied crowd densities with high accuracy, particularly with an MAE of 96.80 and MSE of 139.50 on UCF-CC50, further validates the efficacy of MPD in specialized applications, e.g., crowd counting.

Further, We adapted three augmentation methods tailored for crowd counting and observed MPD outperforms them, as shown in Table 19. Specifically, we compared

Table 17: Comparisons of MPD with other popular augmentation methods. ResNet50: baseline model trained on ImageNet [11] and subsequent models incorporate mentioned augmentation methods. All the models are evaluated on original ImageNet validation set, ImageNet-PD subsets (black background), and Imagenet-E. Both MPD trained models (supervised and SimCLR [7] self-supervised) outperformed on perspective distortion benchmark ImageNet-PD and ImageNet-E while maintain performance on original ImageNet validation set.

Method	Original Val set	ImageNet PD				ImageNet-E
		Top-view (PD-T)	Bottom-view (PD-B)	Left-view (PD-L)	Right-view (PD-R)	
	Top1	Top1/Top5	Top1/Top5	Top1/Top5	Top1/Top5	Top1
ResNet50 [40]	76.13	63.37/83.61	61.15/81.86	65.20/85.13	65.84/85.64	84.28
+Mixup [58]	77.46	65.46/85.15	66.79/85.91	68.02/86.98	68.43/87.22	82.38
+Cutout [12]	77.08	64.27/84.20	62.04/82.86	65.45/85.40	65.61/85.46	79.58
+AugMix [17]	77.53	64.12/84.01	62.90/82.33	65.95/85.94	66.49/86.03	80.97
+Pixmix [18]	77.37	65.52/85.21	64.76/84.54	67.26/86.67	67.56/86.80	82.38
+MPD	76.34	73.00/90.69	72.31/90.81	73.50/91.33	72.91/91.29	85.92
+MPD (SSL)	76.77	73.23/91.39	73.06/91.54	73.55/91.53	73.54/91.65	86.66

Table 18: Ablation on encoder part of crowd-counting method P2P-Net with MPD-CC & MPD-AutoCrowd

Method	Encoder	UCF_CC_50	
		MAE	MSE
P2P-Net	standard VGG16	172.72	256.18
MPD	standard VGG16	101.30	140.65
MPD-AutoCrowd	standard VGG16	96.80	139.50
P2P-Net	standard ResNet50	299.5	382.48
MPD	standard ResNet50	95.10	132.47
MPD-AutoCrowd	standard ResNet50	105.4	124.61

MPD with RandAugment [10], OpenCV PT [59], and AugLy PT [38] across three datasets: SHHA, SHHB, and UCF_CC_50. MPD consistently achieved lower MAE and MSE values, indicating superior performance in handling crowd counting tasks. For instance, MPD-CC achieved a MAE of 51.93 and MSE of 84.3 on SHHA, significantly outperforming AugLy PT’s 56.18 MAE and 107.67 MSE. Furthermore, the samples shown in Fig. 23 illustrates that perspective transform augmentations from [38] and [59] do not capture the non-linearity inherent in real-world perspective distortions, as they primarily rely on affine transformations or translations. This limitation underscores the effectiveness of MPD in capturing complex distortions, leading to its superior performance in crowd counting benchmarks. These results affirm the main paper’s assertion about the potential of MPD in complex visual recognition tasks and demonstrate its versatility in adapting to different encoder configurations while maintaining high performance.

Table 19: Comparisons on augmentation methods in crowd counting. *PT: Perspective Transform.* Metrics: *MAE and MSE errors*

Method	SHHA		SHHB		UCF_CC_50	
	MAE	MSE	MAE	MSE	MAE	MSE
RandAugment [10]	59.81	113.48	10.03	19.84	187.62	271.28
OpenCV PT [59]	57.73	109.82	7.78	13.41	181.05	263.52
AugLy PT [38]	56.18	107.67	8.09	14.05	174.74	259.08
MPD-CC	51.93	84.3	6.73	9.82	101.3	140.65
MPD-AutoCrowd	50.81	85.01	6.61	9.58	96.8	139.5

**Fig. 23:** Applied augmentations on crowd counting

14.2 Transfer learning on Fisheye images

Extended empirical findings, as illustrated in the accompanying graph (Fig. 24), consistently demonstrate that the *ssl:MPD* and *ssl:MPD IB* models excel in multi-label classification tasks on the VOC-360 fisheye dataset, showing a marked improvement in accuracy with an increasing percentage of labels used.

Complementing these results, Tables 20 and 21 provide further insights into the models' performances across a spectrum of MPD probabilities. The data reveal an optimal probability threshold for the MPD transformation application, which maximizes label efficiency. Notably, a probability setting of 0.5 for MPD yields the highest accuracy when all labels are utilized, indicating a non-linear relationship between the transformation probability and classification accuracy.

Table 20: Label efficiency on Fisheye multi-label classification across different probabilities of MPD. ResNet50 is encoder.

Probability	Labels (%)				
	20	40	60	80	100
0.2	38.98	56.52	59.33	82.09	86.76
0.3	37.19	51.50	58.80	78.69	88.03
0.4	38.46	52.46	61.19	78.98	89.24
0.5	38.63	57.14	66.38	76.56	91.63
0.8	38.78	48.45	64.88	81.41	92.69

The results substantiate MPD's versatility and adaptability across different learning approaches and datasets, which is pivotal for real-world applications where perspective distortion is a significant challenge.

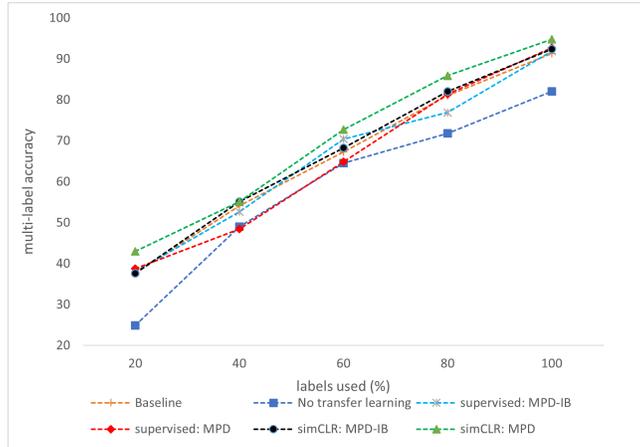


Fig. 24: Label efficiency on Fisheye multi-label classification: Performance comparison across variants of MPD in supervised and self-supervised models with baseline (standard ResNet50) and no-transfer-learning model. All the models evaluated with 20%, 40%, 60%, 80%, and 100% labels.

Table 21: Label efficiency on Fisheye multi-label classification across different probabilities of MPD IB. ResNet50 is encoder.

Probability	Labels (%)				
	20	40	60	80	100
0.2	32.88	51.61	62.62	81.24	87.57
0.3	34.01	51.87	61.35	83.01	89.23
0.4	41.39	48.34	66.68	66.37	86.26
0.5	38.07	52.62	70.44	76.90	91.87
0.8	38.22	48.66	58.65	70.23	89.24

14.3 Person Re-Identification

Our ablation studies further substantiate MPD’s efficacy within the CLIP-ReIdent method by demonstrating additional results on the ResNet50x16 backbone, suggested in the original work [15]. As shown in Table 4, incorporating MPD with a ResNet50x16 encoder significantly outperforms the baseline Clip-ReIdent. The configuration with c_{real} and c_{imag} components ranging from 0.4 to 0.6 achieves the highest mAP of 91.95% without re-ranking and 97.50% with re-ranking. These results are in line with the transformer-based enhancements reported in the main paper and highlight MPD’s broad applicability in addressing perspective distortions across diverse architectures.

14.4 Object Detection

Table 23 presents the object detection performance of MPD-OD on the COCO [30] dataset, with varying probabilities of applying MPD-OD during training. Figure 25

Table 22: Ablation on the person re-identification for MPD (Clip-ReIdent) with CNN backbone: Explores multiple configurations of perceptiveness components with ResNet50x16 encoder. MPD probability set to 0.1, same as before. Min. (component) and Max. (component) defines the range of randomly incorporated perspective distortion. Component selection takes place between c-real and c-imaginary, with equal probability. Direction of each component also chosen with equal probability being positive or negative.

Method	Min. (component)	Max. (component)	mAP (w/o re-ranking)	mAP (with re-ranking)
Clip-ReIdent - ResNet50x16	-	-	88.50	94.90
MPD (Clip-ReIdent) - ResNet50x16	0.2	0.6	91.51	97.40
MPD (Clip-ReIdent) - ResNet50x16	0.2	0.8	90.60	96.38
MPD (Clip-ReIdent) - ResNet50x16	0.4	0.6	91.95	97.50

shows more MPD-OD transformed image examples. The results show that as the probability of applying MPD-OD increases from 0.1 to 0.5, there is a consistent improvement in detection performance for both FasterRCNN [43] and FCOS [47] models. Specifically, FasterRCNN gradually increases IoU=0.50:0.95 from 37.60 to 40.00 and IoU=0.50 from 58.10 to 61.10. Similarly, FCOS shows an improvement from 39.10 to 40.20 in IoU=0.50:0.95 and from 58.50 to 60.30 in IoU=0.50. These trends indicate that higher probabilities of applying MPD-OD during training lead to better object detection performance, highlighting the effectiveness of MPD-OD in enhancing model robustness and accuracy. The consistent performance gains across different probabilities demonstrate the adaptability and efficacy of MPD-OD in handling diverse training scenarios and further strengthen the general usefulness of MPD.

Table 23: Object detection performance of MPD-OD on COCO [30] dataset. Ablations on probability to apply MPD-OD on examples during training. *Metrics: IoU=0.50:0.95 / IoU=0.50*

Method	Original	MPD-OD				
		P=0.1	P=0.2	P=0.3	P=0.4	P=0.5
FasterRCNN [43] (NeurIPS'15)	37.10/55.80	37.60/58.10	39.20/59.40	39.50/60.60	39.70/60.80	40.00/61.10
FCOS [47] (ICCV'19)	38.60/57.40	39.10/58.50	39.50/59.10	39.90/59.70	40.00/59.90	40.20/60.30

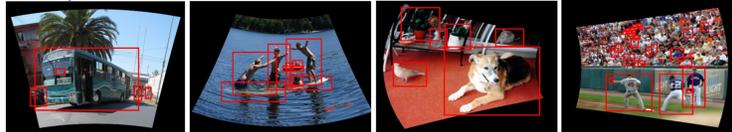


Fig. 25: MPD transformed input & bounding boxes during training