

# Mass Displacement Networks

Natalia Neverova  
nneverova@fb.com

Iasonas Kokkinos  
iasonask@fb.com

Facebook AI Research  
Paris, France

---

## Abstract

Despite the large improvements in performance attained by deep learning in computer vision, one can often further improve results with some additional post-processing that exploits the geometric nature of the underlying task. This commonly involves displacing the posterior distribution of a CNN in a way that makes it more appropriate for the task at hand, e.g. better aligned with local image features, or more compact. In this work we integrate this geometric post-processing within a deep architecture, introducing a differentiable and probabilistically sound counterpart to the common *geometric voting* technique used for evidence accumulation in vision. We refer to the resulting neural models as Mass Displacement Networks (MDNs), and apply them to human pose estimation in two distinct setups: (a) landmark localization, where we collapse a distribution to a point, allowing for precise localization of body keypoints and (b) communication across body parts, where we transfer evidence from one part to the other, allowing for a globally consistent pose estimate. We evaluate on large-scale pose estimation benchmarks, such as MPII Human Pose and COCO datasets, and report systematic improvements.

## 1 Introduction

The advent of deep learning has reduced the amount of hand-engineered processing required for computer vision by integrating many operations such as pooling, normalization, and re-sampling within Convolutional Neural Networks (CNN). The succession of such operations gradually discards the effects of irrelevant signal transformations, allowing the higher layers of CNNs to exhibit increased robustness to small input perturbations. While this invariance is desirable for high-level vision tasks, it can harm tasks such as pose estimation where one aims at precise spatial localization, rather than abstraction.

It is therefore common to apply some form of computer vision-based post-processing on top of CNN-based scores to obtain sharp, localized geometric features. One of the first steps in this direction has been the use of structured prediction on top of semantic segmentation, e.g. by combining image-based DenseCRF [20] inference with CNNs for semantic segmentation [10], training both systems jointly [58], or more recently learning CNN-based pairwise terms in structured prediction modules [9, 23]. All of these works involve coupling decisions so as to reach some consistency in the labeling of global structures, typically coming in the form of smoothness constraints. While this is meaningful for tasks where information is spread out, such as semantic segmentation, we are interested in more general transformations, some of which are illustrated in Fig. 1. For instance, we consider the task of keypoint

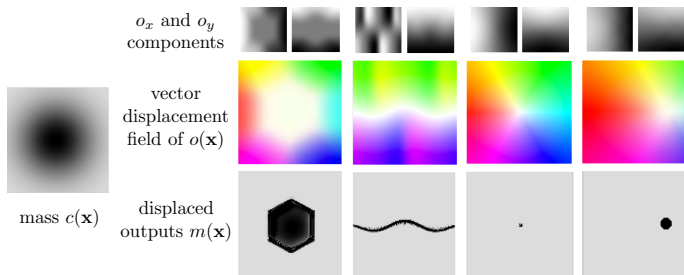


Figure 1: The low spatial resolution of CNNs results in overly smooth heatmaps (*mass*), as shown in the image on the left. Rather than stretch the CNN’s capabilities in order to obtain spatially sharp responses, we propose instead to append a *displacement field* as another CNN output that rearranges the per-pixel confidence scores, lending more evidence to the ground truth positions. The  $\mathbf{o}_x$  and  $\mathbf{o}_y$ -components of different displacement fields  $\mathbf{o}(\mathbf{x})$  are shown on the top (the middle row shows the same components presented as a vector field and displayed in color, for illustrative purposes). These are combined by a Mass Displacement module into a sharp decision, shown in the bottom row.

detection by effectively collapsing the spatially spread output of a CNN into better localized structures. Even though in principle this could be cast in structured prediction terms, the resulting optimization problem amounts to maximizing a submodular function [8] and can only be approximately optimized. We therefore turn to bottom-up, rather than optimization-based methods, and pursue their incorporation in the context of deep learning.

Our starting point is the understanding that requiring high spatial accuracy from a purely CNN-based deep architecture is misusing the network’s abilities: by design, the CNN feature maps get increasingly smooth as we go deeper. We can instead combine these smooth CNN-based classification results with an equally smooth *displacement field* obtained from another CNN branch, indicating to every pixel where its value (or *mass*) should be displaced. This is achieved by separately predicting values of x- and y- components of the displacement vectors of all pixels. Even though the displacement field may be smooth, if its value is accurate, then result can become sharp – in Fig. 1 we are displaying some indicative examples of a smooth response being manipulated by smooth displacement fields that turn it into quite different shapes, that could be appropriate for a variety of visual tasks.

What we are proposing can be understood as *reinventing* geometric voting in the context of deep learning: in a host of computer vision tasks [8, 9, 15, 22, 25, 29] voting can be used to first associate an observation with positions that it supports and then shortlist structures that are supported by multiple observations, e.g. many points voting for a line or a cycle [8], object parts voting for an object’s 2D [22] or 3D pose [30], or many object hypotheses voting for a single object bounding box [16]. Our work was actually motivated by the recent success of such schemes for landmark localization in [27], instance segmentation in [35], and bounding box post-processing in [16].

All of these approaches however are plagued by the heuristic nature of geometric voting, that makes them only applicable as post-processing steps. For example in [27] posterior probabilities are being displaced and then accumulated which results in score maps that can be larger than one – disqualifying them from training with losses appropriate for classification. The authors end up using the cross-entropy loss for the original CNN and the L2 loss

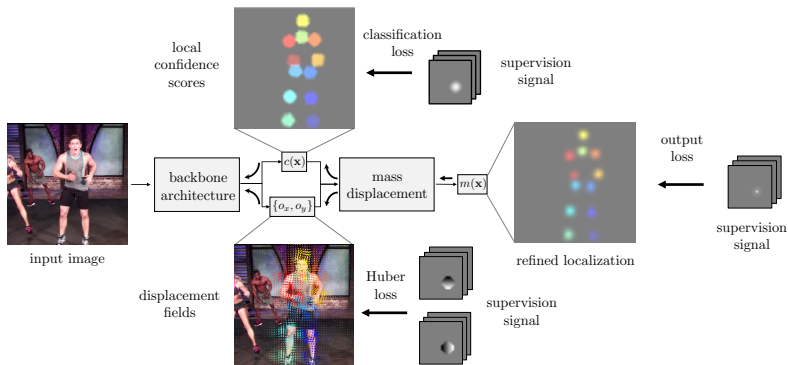


Figure 2: Architecture of a Mass Displacement Network (MDN): the convolutional layers of a CNN are trained with loss functions that allow for some uncertainty in the localization of landmarks, accommodating their inherently smooth responses. A voting operation combines these and collapses the smooth CNN predictions into sharp landmarks. We treat the voting mechanism as a differentiable module and use it for end-to-end training.

for the second stage, while also not training the displacement fields end-to-end – as such it is unclear if the displacement fields are really pointing to the positions that they should. Instead, in this work we develop this somehow ad-hoc post-processing into a module that can easily be combined with existing architectures and trained end-to-end.

In particular, we treat geometric voting as a *differentiable* operation, allowing us to train the CNN-based score maps and displacement fields in an end-to-end manner, ensuring that both arguments to the voting function are optimizing the final system’s performance. Each displaced point is dilated by a kernel to support a region around its novel position, and in the output space every position accumulates evidence from input points that can support it.

Since our approach combines spatial transformation with the geometric manipulation of a probability mass function, we call it a Mass Displacement Network (MDN). Its structure is shown in Fig. 2 and typical outputs in Fig. 3. We explore two tasks: (i) human body landmark localization through *within-part* voting, where the coarse score map of a part is sharpened by a voting process (ii) human pose estimation through *across-part* voting, where every body part score map votes for the presence of other parts. We provide systematic demonstrations of improvements achieved by MDNs over strong baselines on large-scale benchmarks in human pose estimation both in single person and multi-person setups.

Apart from the works mentioned already, our approach has connections to Spatial Transformer Networks (STNs) [20] to bring raw images into correspondence and remove intra-class variation that can be modelled in terms of image deformations. The tacit assumption underlying STNs is that the input and output fields are related by a diffeomorphic transformation, such as a similarity transformation or an affine map, meaning that the dimensionality of structures is preserved. Instead, here we consider transformations that allow us to collapse 2D structures into lower-dimensional structures, such as points, or lines. Furthermore, STNs typically consider a single global parametric transformation, while we have a non-parametric transformation determined by a fully convolutional layer. Finally, as we explain in Sec. 2, STNs are designed like image interpolation operations, and are typically used at the input of a network, while we cater for evidence accumulation, and our module is intended to be




Figure 3: When presented with an image, the three convolutional branches of our network deliver the smooth heatmaps  $c(\mathbf{x})$  and horizontal and vertical displacement fields  $o_x$  and  $o_y$  shown in the middle row (for simplicity, for every kind of output here we display a sum over all planes corresponding to different keypoints). The MDN combines these into the sharper heatmaps  $m(\mathbf{x})$  producing the final joint estimates (on the right).

appended at the end of a network, or generally after some decisions have been produced by a CNN. Finally, in [14] the authors introduced active CNNs, which allow a neuron to pick incoming neurons from input positions determined dynamically through a CNN-based deformation. While this work shares with us the idea of using a CNN-based deformation field, our approach is adapted for collapsing densities and accumulating evidence to certain positions and [14] is tuned for discarding the effect of deformations.

## 2 Mass Displacement Networks

We start by describing the non-probabilistic, geometric voting process currently employed as post-processing in recent works [23, 27] and then propose another principled variant that relies on the noisy-or rule [23] allowing us to use the cross-entropy loss during training. We then turn to the equations used for end-to-end training of the resulting MDN.

**Additive and Noisy-OR voting.** We consider that both our local evidence functions and the output structures reside in a two-dimensional space. In particular we consider that for any position  $\mathbf{x}=(x,y)$  a convolutional network provides us with two outputs: firstly a heatmap, i.e. an estimate of the local confidence for the presence of a feature  $c(\mathbf{x})$ , and secondly an estimate of the predicted structure’s position. The latter is expressed as an horizontal/vertical displacement  $\mathbf{o}(\mathbf{x})$  that should be applied to  $\mathbf{x}$  to obtain the refined estimate  $\mathbf{x}'$ :

$$\mathbf{x}'(\mathbf{x}) = \mathbf{x} + \mathbf{o}(\mathbf{x}) = (x, y) + (o_x(x, y), o_y(x, y)).$$

For landmark localization in within-part voting the displacement field can act like a sharpening correction signal, while for across part voting it reflects supports of relative body part locations. We can accommodate spatial uncertainty in the predicted position by supporting a structure not only at  $\mathbf{x}'$ , but also in the vicinity of the same point. This can be accomplished by dilating the heatmap  $c(\mathbf{x})$  with a kernel, e.g.  $K(\mathbf{x}_o - \mathbf{x}') = \exp(-\|\mathbf{x}_o - \mathbf{x}'\|^2/2\sigma^2)$ , that allows us to smoothly decrease our support as we move further away from  $\mathbf{x}'$ . Combining evidence from multiple points is typically done through summation:

$$m(\mathbf{x}_o) = \sum_{\mathbf{x}} K(\mathbf{x}_o - [\mathbf{x} + \mathbf{o}(\mathbf{x})])c(\mathbf{x}), \quad (1)$$

where for every output position  $\mathbf{x}_o$  we sum the support delivered by all input positions  $\mathbf{x}$ . This has been the setting used for instance in [27] and [23] for landmark localization and

instance segmentation, respectively. In these works a CNN is trained with a cross-entropy loss for  $c(\mathbf{x})$  and a regression loss for  $\mathbf{o}(\mathbf{x})$ , while Eq. 1 is used at test time to deliver more accurate estimates of the desired structures.

The operation in Eq. 1 can be justified in the context of image interpolation, as in the case of Spatial Transformer Networks [20], or in standard Kernel Density Estimation (KDE), but not as a method of accumulating evidence [52]. The main problem, detailed in Appendix A, is that we cannot simultaneously guarantee that the input and output fields both lie in  $[0, 1]$ , so that they can be trained with the cross-entropy loss, and that a confident posterior at  $\mathbf{x}$  will confidently support its displaced replica at  $\mathbf{x} + \mathbf{o}(\mathbf{x})$ , i.e.  $K(\mathbf{0})=1$ .

We can guarantee both requirements by replacing summation with maximization (i.e. perform a ‘‘Transformed-Max-Pooling operation’’). Our experiments with this approach were underwhelming, understandably because we do not accumulate evidence from many points, but rather rely on the single most confident one. Instead we propose differentiable approximations of the maximum operation [0, 28, 52] that allow us to softly combine multiple pieces of evidence while ensuring that the outputs are probabilistically valid.

In particular we use the noisy-or combination rule [28] which provides a probabilistic counterpart to a logical OR-ing operation. We consider that we have  $J$  pieces of evidence about the presence of a feature, each being true with a probability of  $p_j, j=1..J$ . The noisy-or operation expresses the probability  $p$  of the presence of the feature as follows:  $1-p=\prod_{j=1}^J(1-p_j)$ , namely the feature is absent if all supporting pieces of evidence are simultaneously absent – as such, any additional piece of evidence can only increase the estimated value of  $p$ . If now we replace  $p_j$  in the above formula with  $K(\mathbf{x}_o - [\mathbf{x} + \mathbf{o}(\mathbf{x})])c(\mathbf{x})$  we obtain the following rule for combining evidence in the MDN:

$$m(\mathbf{x}_o) = 1 - \prod_{\mathbf{x}} [1 - K(\mathbf{x}_o - [\mathbf{x} + \mathbf{o}(\mathbf{x})])c(\mathbf{x})]. \quad (2)$$

We can use a first-order approximation to obtain Eq. 1 from Eq. 2 if all of the individual terms  $K(\mathbf{x}_o - [\mathbf{x} + \mathbf{o}(\mathbf{x})])c(\mathbf{x})$  are very small, which in hindsight gives some explanation for the practical success of Eq. 1. However, in Eq. 2 we have  $m(\mathbf{x}_o) \in [0, 1]$  which allows us to use the cross-entropy loss throughout training, by virtue of being probabilistically meaningful. Our experiments show that this yields equally good results as the currently broadly used heuristic of regressing to Gaussian functions [0, 26, 50], while being simpler and cleaner.

## 3 Experimental Evaluation

We present experiments in two setups: firstly, *single-person* pose estimation on the MPII Human Pose dataset [11], where the position and scale of a human is considered known in advance. This disentangles the performance of the pose estimation and object detection systems. Secondly, we consider *human pose estimation ‘‘in-the-wild’’* on the COCO dataset [24], where one needs to jointly tackle detection and pose estimation. We use different baselines for both setups, since there is no common strong baseline for both. In both cases MDNs systematically improve strong baselines.

### 3.1 Single person pose estimation

**Dataset & Evaluation:** We evaluate several variants of MDNs on the MPII Human Pose dataset [11] which consists of 25K images containing over 40K people with annotated body joints. We follow the *single person* evaluation protocol, i.e. use a subset of the data with

Model	No voting	Bilinear kernel	$k_f=3$	$k_f=5$	$k_f=7$	$k_f=9$	$k_f=11$	$k_f=13$
Baseline, additive	84.31	87.54	87.70	88.01	88.11	88.15	88.19	88.19
Baseline, noisyOR		87.49	87.63	87.84	87.98	88.08	88.19	88.11
Baseline, max		86.69	86.38	86.22	86.03	85.96	85.34	85.12
Spatial Transformer [10]		88.28	<i>not applicable</i>					
MDN-additive		<b>88.60</b>	×	<b>88.63</b>	×	×	<b>88.61</b>	×
MDN-noisyOR		<b>88.61</b>	×	<b>88.58</b>	×	×	<b>88.32</b>	×

Table 1: Relative performance of the MDN applied to isolated landmarks and trained with different combination rules. All models are based on ResNet-152 and tested on the validation set of MPII Single person. The third baseline is obtained by applying a max operation instead of a sum or a product.  $k_f$  denotes the kernel size in case of a Gaussian kernel.

Model	Resnet-50	Resnet-101	Resnet-152	Hourglass-8
Baseline, no voting	83.29	84.28	84.31	89.24
Baseline, additive, bilinear kernel	86.50	87.50	87.54	89.43
Baseline, noisyOR, bilinear kernel	86.42	87.46	87.49	89.49
MDN-additive, bilinear kernel	87.23	88.42	88.60	<b>89.72</b>
MDN-noisyOR, bilinear kernel	<b>87.25</b>	<b>88.52</b>	<b>88.61</b>	89.64

Table 2: Ablation of interplay between MDN and architectural choices (PCKh on MPII-val).

isolated people assuming their positions and corresponding scales to be known at test time. We follow the standard evaluation procedure of [10] and report performance with the common Percentage Correct Keypoints-w.r.t.-head (PCKh) metric [37]. As in [0, 26], we refine the test joint positions by averaging network predictions obtained with the original and horizontally flipped images.

**Implementation:** We conduct the first exhaustive set of experiments by fine-tuning ImageNet-pretrained ResNet architectures [10]. We substitute the output linear layer and the average pooling that precedes it with a *bottleneck* convolution layer of spatial support  $1 \times 1$  that projects its 2048-dimensional input down to 512 dimensions. This acts like a *buffer layer* between the pretrained network and the pose-specific output layers. As in [27], we reduce the amount of spatial downsampling in such networks by reducing the stride of the first residual module in conv5 block from 2 or 1, and employ atrous convolutions afterwards [10]. As a result, the network takes as an input a cropped image of size  $256 \times 256$ , produces a set of feature planes with spatial resolution of  $16 \times 16$  (rather than  $8 \times 8$ ). These are then bilinearly upsampled to produce the outputs of size  $64 \times 64$ . On top of this common network trunk operate three convolutional branches that deliver the three inputs of the MDN, namely heatmaps ( $\mathbf{x}$ ) and displacement fields  $o_x, o_y$ . Each such branch is a single convolutional layer of spatial support  $1 \times 1$  which maps the 512 feature planes to  $N=16$  dimensions, where  $N$  is the number of landmarks to be localized. The outputs of these branches are passed to the MD layer, which in turn outputs the final refined localizations at the same resolution.

We also present preliminary experimental results with hourglass networks [26], that have even higher performance on MPII – we apply a similar re-purposing as the one outlined above by introducing additional convolutional heads for predicting displacement fields after each stack of the network (where the final estimates for the displacements are obtained by taking a sum over predictions at each step).

**Training:** We test the performance of both additive and noisyOR MDNs. We train the network with three kinds of supervision signals applied to the following outputs:





Figure 4: Examples of MDN-based improvements in human pose estimation through within-part voting (MPII dataset, *Single person track*, ResNet-152) on the left and cross-part voting (COCO dataset, Mask-RCNN-ResNet50-FPN) on the right. Top row: baseline performance, typical failure cases; bottom row: MDN-corrected body pose estimates.

(a) the heatmaps  $c(\mathbf{x})$  learned with pixelwise *binary cross entropy loss*. The supervision signal is presented in the form of binary disks centered at each keypoint location;  
 (b) two displacement planes  $\{o_x, o_y\}$  learned with *robust Huber loss* applied solely in the  $\varepsilon_c$ -vicinity of the ground truth position of every keypoint. The ground truth value for each point  $(x_i^j, y_i^j)$  in the  $\varepsilon_c$ -vicinity of joint  $j$  voting for joint  $k$  is defined as follows:

$$\hat{o}_{x,i}^{(j,k)} = \frac{(x_i^j - x_0^k)}{d} [|x_i^j - x_0^j| \leq \varepsilon_c], \quad (3)$$

where  $(x_0^m, y_0^m)$  is the ground truth position of joint  $m$  and  $d$  is a normalization factor (defined below). The vertical component  $\hat{o}_y$  is defined analogously.

(c) the final refined localizations  $m(\mathbf{x})$ . In this case, depending on the aggregation rule, we apply either *MSE regression loss* (for additive mass displacement) or *binary cross entropy loss* (in case of noisyOR aggregation). The final supervision signal is formulated in the form of a Gaussian (additive MDN) or a binary disk (noisyOR MDN) in the same way as in (a) but with a smaller value of  $\varepsilon_m=1$ . We note that supervising the network with a single loss (c) is possible and produces similar final results but at cost of significantly slower convergence.

All networks are trained using the training set of MPII Single person dataset with artificial data augmentation in the form of flipping, scaling and rotation, as described in [26]. We employ the RMSProp update rule, initial learning rate 0.0025, learning rate decay 0.99, and as in [26] use a validation set of 3k heldout images for our ablation study. We perform evaluation on two separate tasks of within-part and cross-part voting:

(a) **local mass displacement (within body part voting)** – in this setting, the displacement field branches receive their supervision signal in the form of local distributions of horizontal and vertical displacements defined as in 3, where  $j=k$  and  $d=\varepsilon_c$ ;

(b) **global mass displacement (cross body part voting)** – the implementation of the cross voting mechanism is similar to the previous case, but  $j \neq k$  and  $d=X$ , where  $X \times X$  is the output resolution. In this case, we found it more effective to restrict connectivity between joints and perform cross-joint voting along the kinematic tree starting from the center of the body.

**Evaluation results.** In Table 1 we compare Mass Displacement Networks for within-part voting over a set of increasingly complicated baselines: a) a network trained with the

Model	No voting	Bilinear kernel	Gaussian kernel	
			$k_f=5$	$k_f=11$
Baseline, additive	83.96	87.72	87.64	87.73
MDN-additive		<b>88.05</b>	<b>88.08</b>	<b>87.83</b>

Table 3: Relative performance of ResNet-152-MD applied to cross-voting between joints.

Model	Head	Should.	Elbow	Wrist	Hip	Knee	Ankle	Mean-test	Mean-val
Yang et al. [46]	98.5	96.7	92.5	88.7	91.1	88.6	86.0	92.0*	88.5*
Chen et al. [10]	98.6	96.4	92.4	88.6	91.5	88.6	85.7	92.1*	–
Chou et al. [10]	98.2	96.8	92.2	88.0	91.3	89.1	84.9	91.8*	–
Chu et al. [10]	98.5	96.3	91.9	88.1	90.6	88.0	85.0	91.5*	89.4*
Newell et al. [46]	98.2	96.3	91.2	87.1	90.1	87.4	83.6	90.9*	89.4*
Wei et al. [45]	97.8	95.0	88.7	84.0	88.4	82.8	79.4	88.5	–
Insafutdinov et al. [10]	96.8	95.2	89.3	84.4	88.4	83.4	78.0	88.5	–
Belagiannis & Zisserman [9]	97.7	95.0	88.2	83.0	87.9	82.6	78.4	88.1	86.3
Hourglass-8-MDN	98.2	96.4	91.6	87.4	90.8	87.9	84.3	91.3*	89.7*
Resnet-152-MDN	97.7	95.8	90.4	85.1	88.9	85.6	81.6	89.7	88.6

Table 4: Comparison with the state of the art on MPII Single person test set. *Mean-val* denotes PCKh on the validation set. (\*) – models built upon the Hourglass architecture [46].

binary cross-entropy loss with a single objective in the form of a joint heatmap, b) a network outputting the first round of posterior probabilities and displacements independently with following aggregation of corresponding votes in the form of post-processing, i.e. without end-to-end training, c) a modified Spatial Transformer network (STN) [40] aiming on shrinking the produced distributions from iteration to iteration and, just as our architecture, trained end-to-end. In this case, the spatial transformation is not defined globally but instead learned in the form of a vector field describing pixel-wise linear translation. We also evaluate the baseline performance for different filter sizes and gauge the impact of this choice.

We observe that MDNs yield a substantial boost over the different simpler baselines, even when end-to-end training is used, as in the case of STNs (see Fig. 4). The support of the kernel determines the computational complexity of the MD module; we note that by training MDNs end-to-end we achieve excellent results even with  $2 \times 2$  bilinear kernels, rather than using extended Gaussian kernels that make the inference prohibitively slow.

In Table 2 we repeat the same evaluation for different feature extractors with a varying set of network architectures – the results indicate that there is a consistent improvement thanks to the MDN module, and that in all tasks the noisy-or and the additive voting yield virtually identical results. This confirms that we can discard the ad-hoc choice of training the second stage with regression, and replace it with the more meaningful cross-entropy loss.

We next evaluate MDNs on the task of passing information across different joints (*cross-voting*). The corresponding results are shown in Table 3. All models have now been trained to produce three kinds of outputs: posterior probabilities, local displacement fields and across-part displacements. This explains the drop in the baseline’s performance, which was forced to a harder multi-task learning setting (see Table 1 for comparison of the single-task network performance). However, employing an MD layer in the global setting leads to substantial improvement in the localization performance.

Comparison with the state-of-the-art methods is provided in Table 4. It shows that the MDN version of Resnet-152 outperforms all methods not based on Hourglass architecture, while Hourglass-MDN gives a 0.4 point boost over the corresponding baseline and is com-



Method	AP <sup>kp</sup>	AP <sub>50</sub> <sup>kp</sup>	AP <sub>75</sub> <sup>kp</sup>	AP <sub>M</sub> <sup>kp</sup>	AP <sub>L</sub> <sup>kp</sup>	AR <sup>kp</sup>	AR <sub>50</sub> <sup>kp</sup>	AR <sub>75</sub> <sup>kp</sup>	AR <sub>M</sub> <sup>kp</sup>	AR <sub>L</sub> <sup>kp</sup>
Mask R-CNN, keypoints [13]	62.7	87.0	68.4	57.4	71.1	–	–	–	–	–
Mask R-CNN, masks+keypoints [13]	63.1	87.3	68.7	57.8	71.4	–	–	–	–	–
RMPE [9]	61.0	82.9	68.8	57.9	66.5	–	–	–	–	–
CMU-Pose [8]	61.8	84.9	67.5	57.1	68.2	66.5	87.2	71.8	60.6	74.6
G-RMI, COCO only [12]	64.9	85.5	71.3	62.3	70.0	69.7	88.7	75.5	64.4	77.1
Mask R-CNN-MDN, keypoints	63.9	87.2	70.0	58.5	72.3	70.7	91.9	76.2	64.8	78.8

Table 5: Performance of the state-of-the-art pose estimation models trained exclusively on COCO data and tested on COCO `test-dev` (same as in [13]).

petitive with the most complex methods, including, for example, attention mechanisms or adversarial training [0, 11, 12, 13].

Finally, our experiments have shown that stacking several mass displacement modules in different ways (within+across, across+within, or several modules of the same kind) does not further improve performance. This could be explained by the fact that within-part voting is included in cross-joint aggregation (each joint also votes for itself) and, at the same time, dropping cross-joint connections in simple cases allows the model to focus on local aggregation more efficiently. As a result, local voting performs better in the single person setting while cross-joint scheme turned out to be effective in the multi-person scenario.

## 3.2 Multi-person pose estimation

We have obtained similar improvements also on the challenging task of multi-person pose estimation *in the wild*, which includes both people detection and pose estimation. We have built on the recently-introduced Mask-RCNN system of [13] which largely simplifies the task by integrating object detection and pose estimation in an end-to-end trainable architecture.

As in our previous experiments, we have extended the Mask-RCNN architecture with two displacement branches ( $o_x$  and  $o_y$ ) that operate in parallel to the original classification, bounding box regression and pose estimation heads. In the setting of *cross-part* voting, we trained the whole architecture on COCO end-to-end, using identical experimental settings as those reported in [13]. As shown in Table 5, our MDN-based modification of Mask-RCNN yields a substantial boost in performance over the original Mask-RCNN architecture. We also obtain results that are directly comparable to [2], while employing a substantially simpler and faster architecture. Finally, in the supplementary material we show that adding the mass displacement module with additional supervision on displacements further improves performance of detection branches.

## 4 Conclusion

In this work we have introduced Mass Displacement Networks, a principled approach to integrate voting-type operations within deep architectures. MDNs provide us with a method to accumulate evidence from the image domain through an end-to-end learnable operation. We have demonstrated systematic improvements over strong baselines in human pose estimation, in both the single-person and multi-person settings. The geometric accumulation of evidence implemented by MDNs is generic and can apply to other tasks such as surface, curve and landmark estimation in 3D volumetric data in medical imaging, or curve tracking in space and time – we intend to explore these in the future.

## References

- [1] Mykhaylo Andriluka, Leonid Pishchulin, Peter Gehler, and Bernt Schiele. 2d human pose estimation: New benchmark and state of the art analysis. *CVPR*, 2014.
- [2] Boris Babenko. Multiple instance learning: algorithms and applications. Technical report, UCSD, 2008.
- [3] Dana H. Ballard. Generalizing the hough transform to detect arbitrary shapes. *Pattern Recognition*, 13(2):111–122, 1981.
- [4] Olga Barinova, Victor S. Lempitsky, and Pushmeet Kohli. On detection of multiple object instances using hough transforms. *CVPR*, 2010.
- [5] Vasileios Belagiannis and Andrew Zisserman. Recurrent human pose estimation. *FG*, 2017.
- [6] Matthew B. Blaschko. Branch and bound strategies for non-maximal suppression in object detection. *EMMCVPR*, 2011.
- [7] Adrian Bulat and Georgios Tzimiropoulos. Human pose estimation via convolutional part heatmap regression. *ECCV*, 2016.
- [8] Z. Cao, T. Simon, S.-E. Wei, and Y. Sheikh. Realtime multi-person 2d pose estimation using part affinity fields. *CVPR*, 2017.
- [9] Siddhartha Chandra and Iasonas Kokkinos. Fast, exact and multi-scale inference for semantic image segmentation with deep gaussian crfs. *ECCV*, 2016.
- [10] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L. Yuille. Semantic image segmentation with deep convolutional nets and fully connected crfs. *ICLR*, 2015.
- [11] Yu Chen, Chunhua Shen, Xiu-Shen Wei, Lingqiao Liu, and Jian Yang. Adversarial posenet: A structure-aware convolutional network for human pose estimation. *ICCV*, 2017.
- [12] Chia-Jung Chou, Jui-Ting Chien, and Hwann-Tzong Chen. Self adversarial training for human pose estimation. *CVPRW*, 2017.
- [13] Xiao Chu, Wei Yang, Wanli Ouyang, Cheng Ma, Alan L. Yuille, and Xiaogang Wang. Multi-context attention for human pose estimation. *CVPR*, 2017.
- [14] Jifeng Dai, Haozhi Qi, Yuwen Xiong, Yi Li, Guodong Zhang, Han Hu, and Yichen Wei. Deformable convolutional networks. *CoRR*, abs/1703.06211, 2017.
- [15] Juergen Gall, Angela Yao, Nima Razavi, Luc J. Van Gool, and Victor S. Lempitsky. Hough forests for object detection, tracking, and action recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 33(11):2188–2202, 2011.
- [16] Spyros Gidaris and Nikos Komodakis. Object detection via a multi-region & semantic segmentation-aware CNN model. *ICCV*, 2015.

- [17] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. *CVPR*, 2016.
- [18] Kaiming He, Georgia Gkioxari, Piotr Dollar, and Ross Girshick. Mask r-cnn. *CVPR*, 2017.
- [19] Eldar Insafutdinov, Leonid Pishchulin, Bjoern Andres, Mykhaylo Andriluka, and Bernt Schiele. Deepercut: A deeper, stronger, and faster multi-person pose estimation model. *ECCV*, 2016.
- [20] Max Jaderberg, Karen Simonyan, Andrew Zisserman, and Koray Kavukcuoglu. Spatial transformer networks. *NIPS*, 2015.
- [21] Philipp Krähenbühl and Vladlen Koltun. Parameter learning and convergent inference for dense random fields. *ICML*, 2013.
- [22] Bastian Leibe, Ales Leonardis, and Bernt Schiele. Robust object detection with interleaved categorization and segmentation. *International Journal of Computer Vision*, 77 (1-3):259–289, 2008.
- [23] Guosheng Lin, Chunhua Shen, Anton van den Hengel, and Ian D. Reid. Efficient piecewise training of deep structured models for semantic segmentation. *CVPR*, 2016.
- [24] Tsung-Yi Lin, Michael Maire, Serge J. Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. Microsoft COCO: common objects in context. *ECCV*, 2014.
- [25] Subhransu Maji and Jitendra Malik. Object detection using a max-margin hough transform. *CVPR*, 2009.
- [26] Alejandro Newell, Kaiyu Yang, and Jia Deng. Stacked hourglass networks for human pose estimation. *ECCV*, 2016.
- [27] George Papandreou, Tyler Zhu, Nori Kanazawa, Alexander Toshev, Jonathan Tompson, Chris Bregler, and Kevin P. Murphy. Towards accurate multi-person pose estimation in the wild. *CVPR*, 2017.
- [28] J. Pearl. *Probabilistic reasoning in intelligent systems: Networks of plausible inference*. Morgan Kauffman, 1988.
- [29] Nima Razavi, Juergen Gall, Pushmeet Kohli, and Luc J. Van Gool. Latent hough transform for object detection. *ECCV*, 2012.
- [30] Alexander Thomas, Vittorio Ferrari, Bastian Leibe, Tinne Tuytelaars, Bernt Schiele, and Luc J. Van Gool. Towards multi-view object class detection. *CVPR*, 2006.
- [31] Jonathan Tompson, Ross Goroshin, Arjun Jain, Yann LeCun, and Christoph Bregler. Efficient object localization using convolutional networks. *CVPR*, 2015.
- [32] Paul A. Viola, John C. Platt, and Cha Zhang. Multiple instance boosting for object detection. *NIPS*, 2005.
- [33] Shih-En Wei, Varun Ramakrishna, Takeo Kanade, and Yaser Sheikh. Convolutional pose machines. *CVPR*, 2016.

- [34] Christopher K. I. Williams and Moray Allan. On a connection between object localization with a generative template of features and pose-space prediction methods. Technical report, Edinburgh University, 2006.
- [35] Zifeng Wu, Chunhua Shen, and Anton van den Hengel. Bridging category-level and instance-level semantic image segmentation. *CoRR*, abs/1605.06885, 2016. URL <http://arxiv.org/abs/1605.06885>.
- [36] Wei Yang, Shuang Li, Wanli Ouyang, Hongsheng Li, and Xiaogang Wang. Learning feature pyramids for human pose estimation. *ICCV*, 2017.
- [37] Y. Yang and D. Ramanan. Articulated human detection with flexible mixtures of parts. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(12):2878 – 2890, 2013.
- [38] Shuai Zheng, Sadeep Jayasumana, Bernardino Romera-Paredes, Vibhav Vineet, Zhizhong Su, Dalong Du, Chang Huang, and Philip Torr. Conditional random fields as recurrent neural networks. *ICCV*, 2015.