

# Can adversarial networks hallucinate occluded people with a plausible aspect?

Federico Fulgeri, Matteo Fabbri\*, Stefano Alletto, Simone Calderara, Rita Cucchiara

University of Modena and Reggio Emilia, Via P. Vivarelli 10, Modena, 41125, Italy

## ARTICLE INFO

### Keywords:

GAN  
Attribute recognition  
Occlusions

## ABSTRACT

When you see a person in a crowd, occluded by other persons, you miss visual information that can be used to recognize, re-identify or simply classify him or her. You can imagine its appearance given your experience, nothing more. Similarly, AI solutions can try to hallucinate missing information with specific deep learning architectures, suitably trained with people with and without occlusions. The goal of this work is to generate a complete image of a person, given an occluded version in input, that should be a) without occlusion b) similar at pixel level to a completely visible people shape c) capable to conserve similar visual attributes (e.g. male/female) of the original one. For the purpose, we propose a new approach by integrating the state-of-the-art of neural network architectures, namely U-nets and GANs, as well as discriminative attribute classification nets, with an architecture specifically designed to de-occlude people shapes. The network is trained to optimize a Loss function which could take into account the aforementioned objectives. As well we propose two datasets for testing our solution: the first one, occluded RAP, created automatically by occluding real shapes of the RAP dataset created by Li et al. (2016) (which collects also attributes of the people aspect); the second is a large synthetic dataset, AiC, generated in computer graphics with data extracted from the GTA video game, that contains 3D data of occluded objects by construction. Results are impressive and outperform any other previous proposal. This result could be an initial step to many further researches to recognize people and their behavior in an open crowded world.

## 1. Introduction

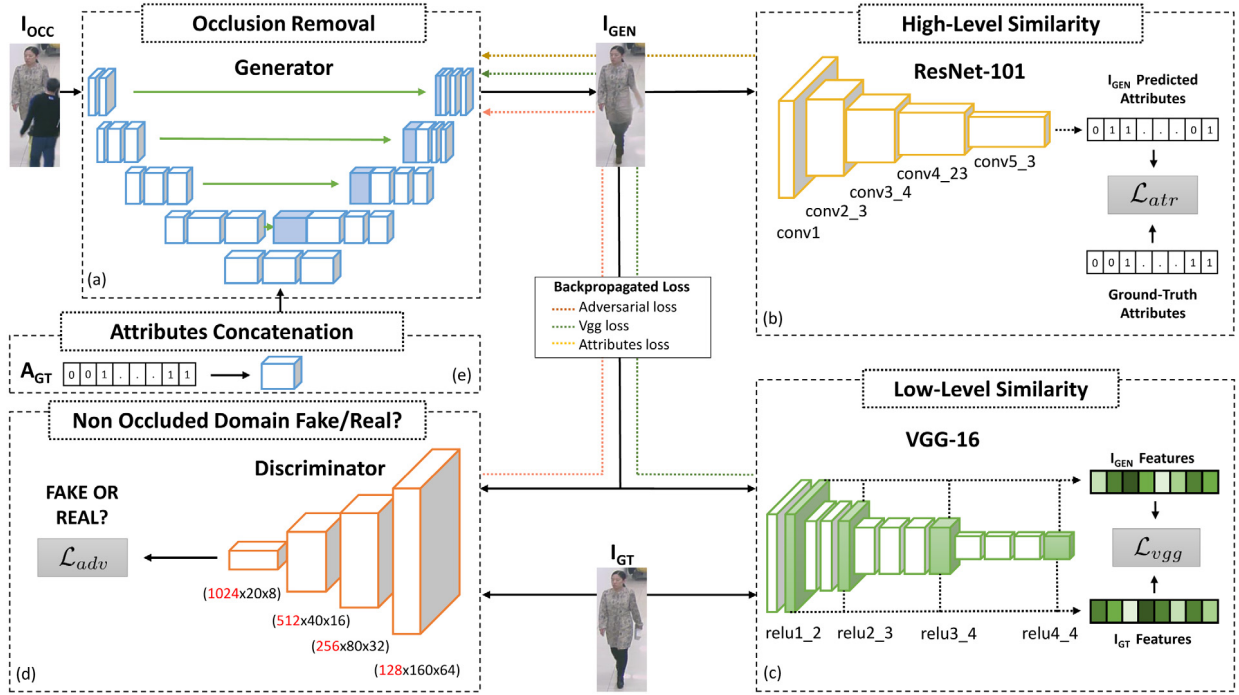
While recent efforts in people detection, recognition, and tracking enabled a plethora of video-surveillance applications, e.g. people identification by Ma et al. (2017), pose estimation by Guler et al. (2018) and action analysis by Herath et al. (2017), occlusion is still an open problem. The occlusion issue is well known in the people detection and tracking literature and generally affects any intelligent video surveillance system, but it is debatable whether a real solution to the problem could exist effectively. In fact, whenever an occlusion occurs we observe a removal of information from the observed scene. The occluded portion of an object, indeed, becomes unknown and, in a Parmenidean sense, it does not exist until it can be observed. For this motivation, most of the literature focused on counteracting the phenomenon conveying occlusion robustness to either detection, tracking, or re-id systems as by Zhuo et al. (2018), Subramaniam et al. (2016), Pan and Hu (2007), Wang et al. (2018a) and Coppi et al. (2016). In the matter of fact, recovering the image content from an occlusion is feasible only in the case where the target has been previously observed e.g. in a video stream. This is the approach

followed also by many tracking solutions which memorize several detected appearance of the person, to discard occlusions as “less frequent accidents” w.r.t. the normal visible appearance. Nevertheless, leveraging on the generative capabilities of GANs by Goodfellow et al. (2014), we aim at demonstrating that it is indeed possible to hallucinate a plausible representation of the occluded content even when it has never been previously observed, i.e. in single images. Following on our previous work on the topic (Fabbri et al., 2017), in this paper, we introduce a novel network that leverages the generative power of GANs for hallucinating the occluded portion of the image without any guidance of an attention mechanism that could provide instance level information about the occluding person. The proposed solution aims at generating or reconstructing the image of a person which could be plausible in many senses: (a) similar to images of real people, observed in the training dataset; (b) acceptable at pixel level as a person shape; (c) capable to preserve similar visual attributes of the ground truth de-occluded image. This is carried out by exploiting solutions for attribute classifications (e.g. male/female, young/old, with/without trousers, etc.) and integrating them in a U-net like generative and adversarial architecture.

\* Corresponding author.

E-mail address: [matteo.fabbri@unimore.it](mailto:matteo.fabbri@unimore.it) (M. Fabbri).

<sup>1</sup> Leveraging on the highly photo-realistic graphics of GTA V video-game.



**Fig. 1.** A schematic representation of the training procedure adopted in our work. The Generator (a) takes the occluded image  $I_{OCC}$  as input and the attributes of the person  $A_{GT}$  (e) as a further conditioning element. The output of the Generator  $I_{GEN}$  is the restored image, with no occlusion. To train the Generator, we fed  $I_{GEN}$  to three different networks: ResNet-101, VGG-16, and the Discriminator. (b) ResNet-101 gives a prediction of the  $I_{GEN}$  attributes which are compared with the ground-truth ones for the  $\mathcal{L}_{atr}$  computation, in order to maximize high-level similarity. (c) The feature maps extracted from different layers of VGG-16 are used to calculate the content loss between  $I_{GEN}$  and  $I_{GT}$  with the aim of encouraging low-level similarity. (d) The Discriminator, which gives the judgment between “real” and “fake” distributions, has to be fooled by the Generator in order to produce images belonging to the non-occluded domain of pedestrians. The Discriminator is trained alongside the Generator to distinguish between generated “fake” images and “real” ones. At evaluation time, only the Generator network and the Attributes Concatenation are used.

Another major problem that arises when dealing with occlusions, through learning-based solutions, is the lack of large-scale datasets providing realistically occluded and non-occluded pairs of images. Most of the proposed solution in literature, like the ones introduced by Fabbri et al. (2017), Ouyang et al. (2016) and op het Veld et al. (2015) paste together different people detections, or manually add random objects or textures to a non-occluded image. These processes ultimately fail to generate realistic data and are thus a liability when employed for training a neural network that aims at resolving the occlusion while keeping the rest of the image coherent (e.g. the background) and preserving the person’s attributes. To address this issue, we propose a novel, fully automatic, way to generate realistic occlusion pairs by exploiting the recent achievements in object segmentation by He et al. (2017). These results are high-fidelity occlusion pairs, where the background of the original image is preserved and the generated occlusion is more realistic. Additionally, thanks to the software provided by Fabbri et al. (2018), we created a massive computer graphics generated dataset,<sup>1</sup> in which we artificially created a large collection of occluded pedestrians. Additionally, we recovered from the game engine their attributes by manually annotating just the models. To our knowledge, this is the first CG dataset for the purpose of de-occluding people having a set of annotated person attributes (e.g. sex, hair color, dress style, etc.).

To summarize, our contributions are threefold:

- We propose a novel generative adversarial network that is able to solve occlusions in pedestrian images by hallucinating the missing parts while keeping both the appearance and the background coherent;
- We devise a new way for synthetically generating occlusion pairs that result in more realistic images when compared to other methods previously employed, also by creating a huge CG dataset;

- We propose a method for conditioning the occluded body part restoration on pedestrian attributes and consequently improving the generation process.

We show by experiments that the design of a conditional GAN that is aware of the attributes can acceptably hallucinate pedestrian and we experimentally demonstrate that this information is helpful in guiding the generation process. Results are interesting in terms of very high accuracy, outperforming other previous methods. We believe that our method could be useful in many computer vision systems, from surveillance, automotive to human behavior understanding tasks.

## 2. Related works

Generative image modeling with deep learning techniques has received lots of attention in recent years. Works on this field can be split into two categories. The first line of works follows the unsupervised setup. Here, the variational autoencoders (VAE) proposed by Rezende et al. (2014) and Kingma and Welling (2013) are the first popular methods which apply a re-parameterization trick to maximize the lower bound of the data likelihood. The most popular methods are indeed generative adversarial networks (GAN) of Goodfellow et al. (2014) and Radford et al. (2015), which simultaneously learn a generator network to generate image samples, and a discriminator network to discriminate generated samples from real ones. GANs are capable of generating sharp images by exploiting the adversarial loss instead of more canonical losses such as L1 or L2.

The second group of works produces images conditioned on either categories, attributes, labels, images or texts. Yan et al. (2016) proposed a Conditional Variational Autoencoder (CVAE) to achieve an image generation conditioned on attributes. On the other hand, Mirza and Osindero (2014) proposed conditional GANs (CGAN) where both the generator and the discriminator are conditioned on extra information to perform category specific image generation. Lassner et al. (2017)

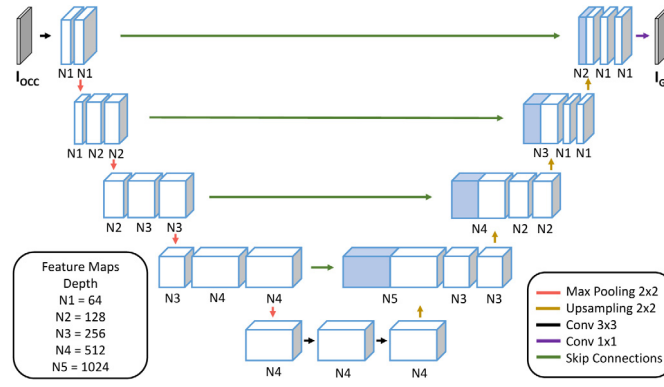


Fig. 2. Architecture of our Generator network with corresponding number of feature maps and kernel sizes. The figure also depicts max-pooling and upsampling operations, along with skip connection gates.

generated people in clothing, by conditioning on the fine-grained body part segments. Reed et al. (2016b) proposed a novel deep architecture and GAN formulation to effectively translating visual concepts from characters to pixels, by adding textual information to both generator and discriminator. They also further investigated the use of additional location, key-points, or segmentation information, to generate images as did by Reed et al. (2016a,c). With only these visual hints as condition and in contrast to our explicit condition on the occluded image, the control exerted over the image generation procedure is still abstract. Many works perform a conditioning over image generation not only on labels or texts but also on images. Zhao et al. (2017) generated multi-view cloth images from only a single view input by proposing a new image generation model that combines the strengths of the variational inference and the GAN framework. Chen and Grauman (2014) tackled the unseen view inference by casting the problem in terms of tensor completion and adopt a factorization approach to accommodate single-view images. Isola et al. (2017) provides a general purpose architecture that is effective at synthesizing photos from label maps, reconstructing objects from edge maps, and colorizing images, among other tasks. Yang et al. (2015), Huang et al. (2017), Yim et al. (2015), Ghodrati et al. (2015) addressed the task of face image generation conditioned on a reference image and a specific face viewpoint. Finally, Yang et al. (2017), Yeh et al. (2017), Pathak et al. (2016), Wang et al. (2018b) tackled the task of image inpainting where large missing regions have to be filled based on the available visual data. Our work can be seen as a particular case of inpainting, where the portion of the image to inpaint is not known a priori.

### 3. Method

The goal of our work is to reconstruct occluded body parts of pedestrians in different surveillance scenarios. Given an image of an occluded pedestrian as the network input, we aim at removing the obstructions and replacing them with body parts that could likely belong to the occluded person. Note that, differently from the task of inpainting, we do not want to guide the network with the information about what portion of the image we want to remove and complete. For this purpose, we want to learn an image transformation between pairs of occluded images  $I_{OCC}$  and not occluded images  $I_{GT}$ .

In order to accomplish this, we propose the training procedure depicted in Fig. 1: our pipeline takes as input the occluded image  $I_{OCC}$ , along with the relative attributes  $A_{GT}$  and outputs the restored image  $I_{GEN}$ .  $I_{GEN}$  is then inputted to the three networks ResNet-101, VGG-16, and the Discriminator in order to compute the three components of our loss. Each loss component is then backpropagated through the input, updating only the Generator's weights.

More precisely, to achieve a full body restoration, we train the Generator network  $G$  as a feed-forward CNN  $G_{\theta_g}$  with parameters  $\theta_g$ .

For  $N$  training pairs images  $(I_{OCC}, I_{GT})$  and their relative attributes  $A_{GT}$ , we solve:

$$\hat{\theta}_g = \arg \min_{\theta_g} \frac{1}{N} \sum_{n=1}^N \mathcal{L}_{total} \left( G_{\theta_g} (I_{OCC}^n, A_{GT}^n), I_{GT}^n \right) \quad (1)$$

Here  $\hat{\theta}_g$  is obtained by minimizing the loss function  $\mathcal{L}_{total}$  described in the next subsection. As a result, our generator network learns a mapping from observed images  $I_{OCC}$  to output image  $I_{GEN}$ . This differs from what did by Isola et al. (2017) and Mirza and Osindero (2014) which use random noise along with the input image.

Following Goodfellow et al. (2014), we further define the Discriminator network  $D_{\theta_d}$  with parameters  $\theta_d$ , that we train alongside  $G_{\theta_g}$  with the aim of solving the adversarial min-max problem:

$$\min_G \max_D \mathbb{E}_{I_{GT} \sim p_{data}(I_{GT})} [\log D(I_{GT})] + \mathbb{E}_{I_{OCC} \sim p_{gen}(I_{OCC})} [\log (1 - D(G(I_{OCC}, A_{GT})))] \quad (2)$$

where  $D(I_{GT})$  is the probability of  $I_{GT}$  being a “real” image while  $(1 - D(G(I_{OCC}, A_{GT})))$  is the probability of  $G(I_{OCC}, A_{GT}) = I_{GEN}$  being a “fake” image. The min-max formulation force  $G$  to fool the  $D$ , which is adversarially trained to distinguish between generated “fake” images and “real” ones. Thanks to this approach, we obtain a Generator network capable of learning solutions that are similar to not occluded images thus indistinguishable by the Discriminator. Note also that, differently from what did by Isola et al. (2017), we do not concatenate input images  $I_{OCC}$  to the “fake” images  $I_{GEN}$  or to the “real” images  $I_{GT}$  as Discriminator input.

#### 3.1. Loss function

The definition of the loss function  $\mathcal{L}_{total}$  is crucial for the effectiveness of our Generator network. We propose the following loss formulation, composed by a weighted combination of three components:

$$\mathcal{L}_{total} = \underbrace{\mathcal{L}_{adv}}_{\text{adver. loss}} + \underbrace{\lambda_1 \cdot \mathcal{L}_{vgg}}_{\text{cont. loss}} + \underbrace{\lambda_2 \cdot \mathcal{L}_{atr}}_{\text{attr. loss}} \quad (3)$$

The intuition behind this loss formulation is that we want the generated images to contain real people (thanks to  $\mathcal{L}_{adv}$ ), to have similar feature representations (thanks to  $\mathcal{L}_{vgg}$ ) and to preserve visual attributes (thanks to  $\mathcal{L}_{atr}$ ) w.r.t. their non occluded ground truth versions.

The first term of Eq. (3) is the adversarial loss  $\mathcal{L}_{adv}$ . This component encourages the Generator network  $G$  to generate images belonging to the not occluded domain of pedestrians by fooling the Discriminator network  $D$ . The adversarial component relative to the Generator network is calculated as follows:

$$\mathcal{L}_{adv} = - \sum_{n=1}^N \log D(G(I_{OCC}, A_{GT})) \quad (4)$$



**Table 1**


Classification performances of our ResNet-101 on RAP dataset.

Method	mA	Accuracy	Precision	Recall	F1
ACN <a href="#">Sudowe et al. (2015)</a>	69.66	62.61	<b>80.12</b>	72.26	75.98
DeepMAR <a href="#">Li et al. (2015)</a>	73.79	62.02	74.92	76.21	75.56
DeepMAR* <a href="#">Li et al. (2016)</a>	74.44	63.67	76.53	77.47	77.00
HP-Net <a href="#">Liu et al. (2017)</a>	76.12	65.39	77.33	78.79	78.05
ACN-Res50 <a href="#">Fabbri et al. (2017)</a>	<b>79.73</b>	64.13	76.96	78.72	77.83
Ours	78.66	<b>66.23</b>	77.85	<b>79.71</b>	<b>78.77</b>

**Table 2**

Detailed comparison between various pedestrian attribute datasets.

Dataset	# Scenes	# Samples	# Attributes	Min. resolution	Max. resolution
PETA <a href="#">Deng et al. (2014)</a>	–	19,000	61 (+4)	17 × 39	169 × 365
RAP <a href="#">Li et al. (2016)</a>	26	41,585	69 (+3)	36 × 92	344 × 554
PA-100K <a href="#">Liu et al. (2017)</a>	598	100,000	26	50 × 100	758 × 454
AiC	512	125,000	24	35 × 85	602 × 1080

**Fig. 3.** Examples from the AiC dataset exhibiting its variety in viewpoints, illuminations and scenarios.

Where  $D(G(I_{OCC}, A_{GT}))$  is the probability that  $G(I_{OCC}, A_{GT})$  is classified as “real” by the discriminator network. Minimizing  $\log D(G(I_{OCC}, A_{GT}))$  instead of  $\log D[1 - (G(I_{OCC}, A_{GT}))]$  is preferred in order to reach a better gradient behavior as indicated by [Goodfellow et al. \(2014\)](#). As a possible drawback, the images produced by the Generator network  $G$  are forced to be realistic thanks to the Discriminator network  $D$ , but they can be unrelated to the original input. For instance, the output could be a plausible image of a pedestrian displaying a very different aspect w.r.t. the input image. Thus, is essential to mix the adversarial loss  $\mathcal{L}_{adv}$  with an additional loss, such as L1 or L2, that evaluate the per-pixel distance between the generated and the ground truth image. Usually, training a network using such losses leads to reasonable solutions. However, the outputs appear blurred with lack of high-frequency details.

A possible solution for generating sharper results is to adopt a different content loss, like the perceptual loss introduced by [Johnson et al. \(2016\)](#) and used also by [Ledig et al. \(2017\)](#) and [Kupyn et al. \(2017\)](#):

$$\mathcal{L}_{vgg(i,j)} = \frac{1}{W_{i,j} H_{i,j}} \sum_{x=1}^{W_{i,j}} \sum_{y=1}^{H_{i,j}} (\phi_{i,j}(I_{GT})_{x,y} - \phi_{i,j}(I_{GEN})_{x,y})^2 \quad (5)$$

where  $W_{i,j}$  and  $H_{i,j}$  are the dimensions of the feature maps  $\phi_{i,j}$  obtained by the  $j$ th convolution after ReLU activation and before the  $i$ th max-pooling layer within the VGG16 network, pre-trained on ImageNet by [Deng et al. \(2009\)](#), as done by [Johnson et al. \(2016\)](#).

The  $\mathcal{L}_{vgg}$  that we employed in our work is based on the sum of different intermediate layers of VGG16:

$$\mathcal{L}_{vgg} = \sum_{i,j \in L} \mathcal{L}_{vgg(i,j)} \quad (6)$$

where  $\mathcal{L}_{vgg(i,j)}$  is taken from Eq. (5) and  $L$  is the set of used activations. Rather than encouraging the pixels of the output image  $I_{GEN}$  to exactly match the pixels of the target image  $I_{GT}$ , we instead encourage them to have similar feature representations as computed by the VGG16 network. As demonstrated by [Johnson et al. \(2016\)](#) and [Mahendran and Vedaldi \(2015\)](#), minimizing the content loss for higher layers do not preserve color and textures. As we reconstruct from early layers, instead, images tend to be perceptually similar to the target image  $I_{GT}$  in terms of color and texture.

Since our main purpose is not limited to naively restore the occluded parts of pedestrians, but also to maintain and highlight their attributes, we introduced an additional loss component  $\mathcal{L}_{att}$ . Like for the perceptual loss  $\mathcal{L}_{vgg}$ , we used a classification network as loss function. In

particular, we adapted ResNet-101 by He et al. (2016), pre-trained on ImageNet, to the task of multi-attribute classification. More precisely, we replaced the last two layers (the average pooling and the last fully connected layer) in order to fit the desired input shapes. Differently, from the VGG loss, we work on a higher level of abstraction, forcing the Generator network to produce images that exhibit characteristics coherent with the attributes of the person. In this case, we used a weighted binary cross entropy:

$$\mathcal{L}_{atr} = - \sum_{i=1}^{N_A} \exp(1 - r_i) \cdot (y_i \cdot \log(\psi_i(I_{GEN}))) + \exp(r_i) \cdot (1 - y_i) \cdot \log(1 - \psi_i(I_{GEN})). \quad (7)$$

Here,  $N_A$  is the number of attributes classified by the ResNet-101,  $r_i$  is the positive ratio of  $i$ th attribute.  $\psi$  is the output of our attribute classification network and  $y_i$  is the  $i$ th ground truth label.

### 3.2. Networks architecture

**Generator network.** Our Generator structure differs from those presented by Radford et al. (2015) and Fabbri et al. (2017): following Ronneberger et al. (2015) and Isola et al. (2017) we propose the “U-Net” like architecture depicted in Fig. 2. In particular, the structure of our network slightly differs from the one described by Ronneberger et al. (2015) and Isola et al. (2017). The network is composed by 4 down-sampling blocks and a specular number of up-sampling components. Each down-sampling block consists of 2 convolutional layers with a  $3 \times 3$  kernel. Each convolutional layer is followed by a batch normalization and a ReLU activation. Finally, each block has a max-pooling layer with stride 2. The up-sampling part has a very similar but overturned structure, where each block is composed by an up-sampling layer of stride 2. After that, each block is equipped with 2 convolutional layers with a  $3 \times 3$  kernel. The last block has an additional  $1 \times 1$  kernel convolution which is employed to reach the desired number of channels: 3 RGB channels in our case. A *tanh* has been used as final activation. We additionally inserted skip connections between mirrored layers, in the down-sampling and up-sampling streams, in order to shuttle low-level information between input and output directly across the network. Eventually, padding is added to avoid cropping the feature maps coming from the skip connections and concatenate them directly to the up-sampling blocks outputs. Roughly speaking, our task can be seen as a particular case of image-to-image translation, where a mapping is performed between the input image and the output image. Additionally, for the specific problem we are considering, input and output share the same underlying structure despite differing in superficial appearance. Therefore, a rough alignment is present between the two images. In fact, all the non-occluded parts that are visible in the input images must be transferred to the output with no alterations. The structure of the U-Net lends itself optimally to our task, and the skip connections are fundamental for the conservation of the non-occluded image content. In this way, useful low-level information is not lost during the encoding passage: by leveraging this kind of information, we are able to maintain the appearance of visible parts in the image.

**Discriminator network.** The Discriminator, instead, aims to determine if an image is true or if it has been generated. In particular, the structure is similar to the one proposed by Radford et al. (2015), composed by 4 convolutional layers with kernel size  $5 \times 5$ . The resulting features are followed by one sigmoid activation function in order to obtain a probability for the classification problem. We use batch normalization before every Leaky ReLU activation, except for the first layer.

### 3.3. Training details

We trained our GAN with  $320 \times 128$  resized input images while simultaneously providing the target image in order to compute the supervised losses. We adopted the standard approach by Goodfellow et al. (2014) to optimize the networks alternating gradient descent updates between the Generator and the Discriminator with  $K = 1$ . Data augmentation is performed by randomly flipping the images horizontally. We used mini-batch SGD applying the Adam solver with momentum parameters  $\beta_1 = 0.5$  and  $\beta_2 = 0.999$ , learning rate  $2 \cdot 10^{-4}$  and a batch size of 20. Each training is performed using a Titan Xp GPU.

## 4. Datasets

We evaluated our method on the RAP dataset, proposed by Li et al. (2016), comparing state-of-the-art methods and performing the ablation study over each loss employed. In addition, we further propose a new large-scale computer-graphics dataset AiC for pedestrian attribute recognition in crowded scenes. Differently, from existing publicly available datasets, AiC is mainly focused on occlusion events.

### 4.1. RAP dataset

RAP by Li et al. (2016) is a very richly annotated dataset with 41,585 pedestrian samples, each of which is labeled with 72 attributes as well as viewpoints, occlusions, and body parts information. In order to evaluate our method, we corrupted the dataset with occlusions. Differently, from what did by Fabbri et al. (2017), where obstructions were created by cutting parts of images according to regular geometric shapes, we have adopted a more sophisticated approach that has led us to more realistic results. By exploiting the state-of-the-art performances of Mask R-CNN proposed by He et al. (2017), pre-trained on the COCO Dataset (Lin et al., 2014), we produced segmentation masks for each person in the RAP dataset. The computed silhouettes were then used to crop people’s shapes from the dataset. Those figures are then used to reproduce the occlusions, simply by randomly overlapping the crops to each image sample of RAP dataset. In addition, to reduce the visual gap between the original image and the overlapped person, we performed a Gaussian blurring. However, this is not applied to the whole image but only to the area given by the difference between an expansion and an erosion of the segmentation mask of the overlapping image. The only constraint that we have introduced is that the occluding person must not occupy the portion of the image that has the y coordinate that exceeds the 6/7 of the image height. Each sample is computed as follows:

$$I_{OCC} = I_{GT^1} \odot \neg \alpha(\beta(I_{GT^2})) + \alpha(\beta(I_{GT^2}) \odot I_{GT^2}) \quad (8)$$

where  $\beta(I_{GT^2})$  is the binary mask generated using Mask R-CNN and morphology operations and  $\alpha$  is a function used to translate the overlap section randomly over the destination image  $I_{GT^1}$ . The dataset is already organized in 5 random splits. Each of which contains 33,268 images for training and 8317 for testing. As did by Li et al. (2016), due to the unbalanced distribution of attributes in RAP, we selected the 51 attributes that have the positive example ratio in the dataset higher than 0.01.

### 4.2. Aic dataset

Most of the publicly available pedestrian attribute datasets, like RAP by Li et al. (2016), PETA by Deng et al. (2014) and PA-100K by Liu et al. (2017) does not contemplate occlusion events (see Table 2 for further details). They only provide samples of full visible people, completely ignoring crowded situations of pedestrians occluding each other (which is indeed common in urban scenarios). To overcome this limitation, we propose the Attributes in Crowd dataset, a novel

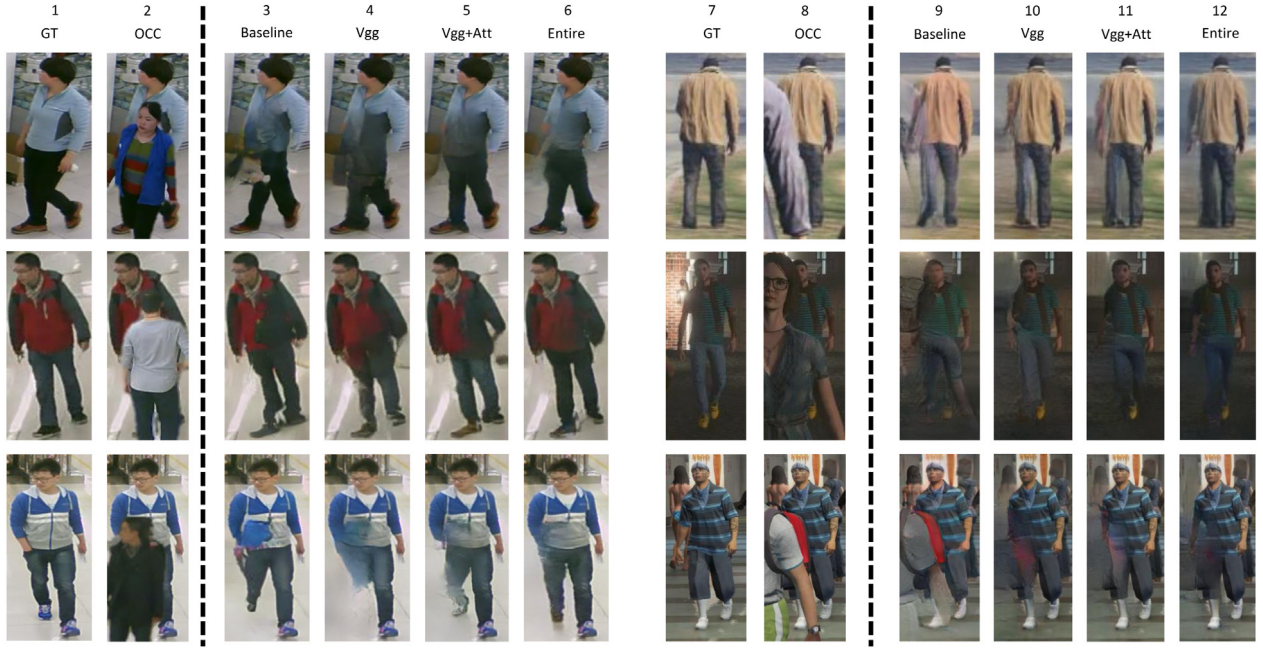


Fig. 4. Qualitative results based on the ablation study on RAP dataset (leftmost) and AiC dataset (rightmost). GT columns indicate ground truth images while in the OCC columns are presented the input occluded images. Columns 3 and 9 indicate the outputs of our baseline, where adversarial loss and MSE are used. Columns 4 and 10 represents results of the VGG loss. On 5 and 11 we have results of experiments using all the 3 losses combined: adversarial loss, VGG loss, and attribute loss. Finally, columns 6 and 12 show results where attributes are injected as input to the network.

Table 3

Ablation study results on RAP dataset.

Method	mean Accuracy	Accuracy	Precision	Recall	F1	SSIM	PSNR
Occlusion	65.74	51.06	68.72	64.36	66.47	0.7153	14.57
Baseline	70.74	56.55	70.61	71.78	71.19	0.7982	20.31
VGG loss	72.48	58.89	72.58	73.56	73.06	<b>0.8293</b>	<b>20.88</b>
VGG and attr. loss	72.18	59.59	73.51	73.72	73.62	0.8239	20.65
VGG and attr. loss (+input attr.)	<b>81.1</b>	<b>74.8</b>	<b>84.29</b>	<b>85.61</b>	<b>84.94</b>	0.8274	20.7
GT data	78.66	66.23	77.85	79.71	78.77	–	–

Table 4

Ablation study results on AiC dataset.

Method	mean Accuracy	Accuracy	Precision	Recall	F1	SSIM	PSNR
Occlusion	72.24	45.77	48.78	79.03	60.32	0.6148	18.38
Baseline	72.72	45.48	48.23	80.87	60.42	0.6236	20.49
VGG loss	78.12	53.11	55.52	85.65	67.37	0.7088	21.5
VGG and attr. loss	78.37	53.3	55.73	85.46	67.46	<b>0.7101</b>	<b>21.81</b>
VGG and attr. loss (+input attr.)	<b>90.86</b>	<b>72.15</b>	<b>74.0</b>	<b>95.1</b>	<b>83.23</b>	0.6986	21.47
GT data	91.89	74.87	76.80	95.43	85.11	–	–

synthetic dataset for people attribute recognition in presence of strong occlusions. AiC features 125,000 samples, all being a unique person, each of which is automatically labeled with information concerning sex, age etc. The dataset is split into 100,000 samples for training and 25,000 for testing purposes. Each of the 24 attributes is present at least in a 10% of samples which highlight a good balance in terms of labels. The collected samples feature a vast number of different body poses, in several urban scenarios with varying illumination conditions and viewpoints. Skeleton joints are also available for each identity. Joints are additionally labeled with an occlusion flag which tells if the specific body part is directly visible from the camera point of view. Moreover, each image sample has his vanilla version where each obstacle is removed from the image. Thus, for each occluded pedestrian, we know exactly how it really is behind the occlusion (this is obviously not achievable in real environments). Fig. 3 exhibits some examples of the dataset. To foster future research on this topic, the dataset will

be publicly released upon publication. AiC was created by exploiting the highly photo-realistic video game *Grand Theft Auto V* developed by Rockstar North.

## 5. Experimental results

In this section, we provide details about the metrics adopted, followed by a detailed ablation study that presents qualitative and quantitative results for three different combinations of losses (that we added to the adversarial loss): MSE loss, VGG loss and a combination of VGG loss and attribute loss. We also investigate how the information about the attributes of a person can enhance the quality of the produced images. Additionally, we explain the choice of different hyperparameters, exploring their impacts. Finally, we compare our method with the most related works presented by Isola et al. (2017) and Fabbri et al. (2017).





Fig. 5. Qualitative results on both RAP and AiC datasets. (first line) an example using a configuration of  $\lambda_1 = 0$  and  $\lambda_2 = 15$  on RAP: the color of the jacket mutates from pink to gray to facilitate the classification, as the majority of jackets in the dataset are dark. (second line) an example using a configuration of  $\lambda_1 = 10$  and  $\lambda_2 = 0.1$  on AiC: blurring the occlusion and not hallucinating new body parts results in a better strategy to facilitate ResNet-101. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

### 5.1. Evaluation metrics

Evaluating the quality of synthesized images is an open and challenging problem as stated by Salimans et al. (2016). Traditional metrics

such as per-pixel MSE do not estimate joint statistics of the result, and therefore do not extrapolate the full structure of the result. In order to more holistically evaluate the visual quality of our results, we employed two tactics. Firstly, we compared the performance of the proposed model through metrics directly calculated over the reconstructed images. Specifically, we adopted the structural similarity SSIM and the peak signal-to-noise ratio PSNR. Secondly, we measured the capability of the proposed network of being able to preserve original attributes, like gender, hairstyle or wearing jacket, by exploiting the ResNet-101 network of He et al. (2016) trained on the task of multi-attribute classification. Thus, following Li et al. (2016), Fabbri et al. (2017) and Liu et al. (2017), we provide five evaluation metrics for the attribute classification task, namely mean Accuracy, Accuracy, Precision, Recall and F1.

**ResNet-101 classification network.** We trained the network with  $320 \times 128$  resized images with Adam as optimizer and learning rate set to  $2 \cdot 10^{-4}$ . In Table 1 a comparison on the classification task with other state-of-the-art networks on RAP dataset is presented. The same network is trained independently for each dataset, in order to provide reliable metrics for both RAP and AiC.

### 5.2. Ablation study

As previously stated, we investigated three loss combinations in order to clarify and highlight the solutions adopted in our work:

- **Baseline:** the Baseline architecture uses, in conjunction with the adversarial loss, the MSE loss as content loss;
- **VGG loss:** differently from the Baseline, we replaced the MSE loss with the VGG loss. The layers (1, 2), (2, 2), (3, 3) and (4, 3) are chosen as the set  $L$  of activations on Eq. (6). In Eq. (3), we set  $\lambda_1$  to 10 and  $\lambda_2$  to 0 (further details about  $\lambda_1$  and  $\lambda_2$  are presented in the next subsection);
- **VGG loss + Attr. loss:** in this case, all the three losses are employed. The VGG loss always refers to the same four activation layers. The Attribute loss is computed between the output of the ResNet-101

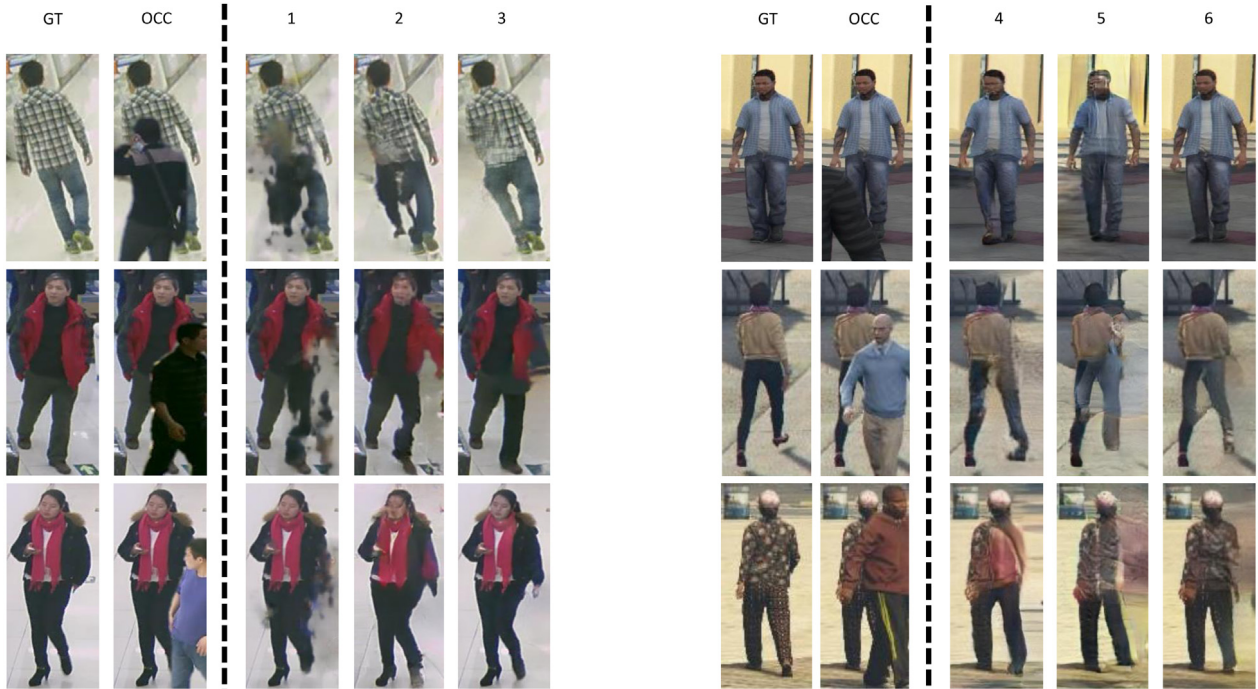


Fig. 6. Qualitative comparison with state-of-the-art approaches: results are presented for both RAP (leftmost) and AiC (rightmost). GT columns indicate ground truth images while in the OCC columns are presented the input occluded images. Columns 1 and 4 are the images recovered by Pix2Pix by Isola et al. (2017). On 2 and 5 are presented results obtained from the method used by Fabbri et al. (2017). The last two columns, 3 and 6, show our best comparable approach output (Vgg loss + Attr. loss).

classification network computed on the generated images and the ground truth labels provided by the datasets. In Eq. (3), we set  $\lambda_1$  to 10 and  $\lambda_2$  to 5 for RAP and to 0.01 for AiC. Note that we did not use all the available attributes of RAP dataset, but only the first 51 for the reason explained at the end of Section 4.1. For AiC dataset, instead, we used all the available attributes.

In order to further investigate how some additional information about the attributes can improve the restoration process, we performed a further experiment where attributes are fed as input to the network, along with the occluded image:

- *Entire*: in this setup, we adopted both the VGG loss and the Attribute loss, along with the adversarial loss. Differently, from our main method, attributes are injected directly to the main flow of the Generator network. Specifically, the attribute vector of the occluded pedestrian is fed to a fully connected layer in order to produce a feature vector that is reshaped to match the bottleneck dimension of our Generator network.

Fig. 4 shows some qualitative results. The baseline performs considerably worse than the other setups, not being able to completely remove the occlusions on AiC (column 9 of Fig. 4). This is probably due to the fact that AiC is a more challenging dataset compared to our corrupted version of RAP. For the same reason, RAP results are overall more appealing than the ones of AiC. Moreover, no substantial difference appears between the other setups, highlighting the fact that the VGG loss is the main component that guides the network to produce high-quality results.

Tables 3 and 4 present quantitative results for RAP and AiC respectively based on our ablation study. The tables also provide metrics referred to the occluded images before the restoration process. By observing the tables, we can state that, despite being visually indistinguishable, the images obtained from the VGG loss and from our Entire configuration produce very different results in terms of attribute metrics. We can also observe that there is no substantial difference between the VGG loss and the VGG loss with Attributes loss. In fact, RAP shows a gap of one percentage point in almost all the classification metrics, while AiC shows very little differences, probably due to the more challenging nature of AiC. Moreover, Table 3 shows that the Entire setup reach higher scores compared to the upper bound of the ground truth images. Also Table 4 shows performances that are close to the ground truth metrics when we input attribute information directly to the Generator. In fact, with attributes as input, the Generator network, by restoring the occluded images, is able to produce an output that has enhanced attribute characteristics (although this is not visible to the naked eye). As can be shown in the next subsection, further forcing the generation output on classification metrics, we can reach results that exceed the ground truth upper bound even on AiC, at a price of a drop on reconstruction metrics.

### 5.3. Hyperparameter optimization

Hyperparameter tuning is a crucial aspect in designing machine learning frameworks, as the performance of an algorithm can be highly dependent on the choice of hyperparameters. In fact,  $\lambda_1$  and  $\lambda_2$  were selected using a grid search technique. In particular, we searched for a trade off between classification metrics (accuracy, precision, recall, f1) and pixel-level reconstruction metrics (PSNR, SSIM). We performed a different grid search for four different configurations combining each dataset with the two main setups: VGG loss + Attr. loss and the Entire pipeline.

For what concerns the VGG loss + Attr. loss setup, we observed that, in general, a configuration with  $\lambda_1 \gg \lambda_2$  brings to better pixel-level reconstruction metrics but poor classification performances. On the other hand, solutions with  $\lambda_1 \ll \lambda_2$  show good classification performances but low pixel-level reconstruction metrics. Also, increasing  $\lambda_1$  over the

value of 10 does not further improve PSNR and SSIM metrics (for both RAP and AiC). The same behavior happens for  $\lambda_2$ : the classification metrics do not improve for values greater than 5 (for RAP) and 0.01 (for AiC). This difference of  $\lambda_2$  between the two datasets may be caused by the fact that AiC is more challenging than RAP. In fact, during training, the Attributes Loss on AiC is orders of magnitude greater than the same loss on RAP, thus, a smaller  $\lambda_2$  is needed to maintain the balance between the losses.

For what concern the Entire pipeline, we observed a different behavior on  $\lambda_2$ : increasing  $\lambda_2$  does steadily improve the classification metrics (reaching up to 98.89 mean Accuracy with  $\lambda_2 = 5$ ) while drastically decreasing PSNR and SSIM. This behavior happens on both RAP and AiC. By giving more importance to the Attributes Loss, the Generator network is able to enhance attribute characteristics to the point that they are highly recognizable by the classification network, at the price of not maintaining low-level similarity. Fig. 5 shows a direct consequence at qualitative level on both RAP and AiC. The first line depicts an extreme configuration of  $\lambda_1 = 0$  and  $\lambda_2 = 15$  on RAP. With no low-level constraints, the Generator network is able to mutate the color of the jacket to facilitate the ResNet-101 “jacket attribute” recognition. The second line of Fig. 5, instead, shows an example obtained using  $\lambda_1 = 10$  and  $\lambda_2 = 0.1$  on AiC. In this case, the behavior is completely different: due to the high diversity of attributes in AiC, the Generator learns to simply remove the obstacle, not adding (hallucinating) many details to the removed portion of the image. Adding imprecise details would, in fact, mislead the attribute classification network.

### 5.4. Comparison against state-of-the-art techniques

Since our task of de-occlusion is novel, there are no direct works to compare with. So, to match the results of our network, in addition to our previous work, we also retrained the Pix2Pix framework on both RAP and AiC.

*Our previous work.* Like our current method, Fabbri et al. (2017) exploits an adversarial based framework to achieve a translation from an occluded-pedestrian domain to a completely visible body domain. The main difference with our current method resides in the loss formulation: Fabbri et al. (2017) minimizes a combination of adversarial loss and sum of squared error loss (SSE), completely ignoring high-level and low-level similarities. Another important difference lies in the Generator architecture: our previous work uses a simple hourglass architecture with no *skip connections*, while in our current method we adopted a U-net based solution. The U-net architecture shows better performances in tasks where some input information has to be shuttled directly to the output with no variation. In fact, as can be seen in Fig. 6, our previous work fails to preserve the portions of the image that should remain unchanged (especially the faces).

*Pix2Pix.* Isola et al. (2017) investigates conditional adversarial networks as a general-purpose solution to image-to-image translation problems. As in our Generator network, Pix2Pix exploits a U-net based architecture. The only substantial architectural difference is in the number of convolutional layers before each downsampling and after each upsampling operation. Also, the Discriminators differs: Pix2Pix uses a patch level discriminator that only penalizes structure at the scale of patches, while in our work we adopt an image level discriminator that takes the whole image as input. A patch level discriminator models the image as a Markov random field, assuming independence between pixels separated by more than a patch diameter. This is indeed not the case when dealing with images of people. In fact, for example, the skin color of the face should match the skin color of the hands. Also, the trousers are usually made of the same color. Another significant difference lies in the content loss: Pix2Pix, like our previous work, uses a pixel-level loss (L1 instead of SSE), assuming pixel independence, and forcing pixels of the output image to exactly match the pixels of the target image. In our work, instead, we exploit a combination of



**Table 5**

Comparison with the state-of-the-art method on RAP dataset.

Method	mA	Accuracy	Precision	Recall	F1	SSIM	PSNR
Occlusion	65.74	51.06	68.72	64.36	66.47	0.7153	14.57
Pix2Pix <a href="#">Isola et al. (2017)</a>	69.49	52.05	65.07	70.06	67.47	0.7348	17.91
<a href="#">Fabbri et al. (2017)</a>	65.92	51.44	65.77	67.94	66.84	0.6798	18.4
Ours	<b>72.18</b>	<b>59.59</b>	<b>73.51</b>	<b>73.72</b>	<b>73.62</b>	<b>0.8239</b>	<b>20.65</b>

**Table 6**

Comparison with the state-of-the-art method on AiC dataset.

Method	mA	Accuracy	Precision	Recall	F1	SSIM	PSNR
Occlusion	72.24	45.77	48.78	79.03	60.32	0.6148	18.38
Pix2Pix <a href="#">Isola et al. (2017)</a>	71.93	44.27	46.75	81.61	59.45	0.6351	21.22
<a href="#">Fabbri et al. (2017)</a>	67.14	38.21	40.61	79.9	53.85	0.573	20.11
Ours	<b>78.37</b>	<b>53.3</b>	<b>55.73</b>	<b>85.46</b>	<b>67.46</b>	<b>0.7101</b>	<b>21.81</b>

high-level and low-level consistency by encouraging the overall images to have similar feature representations as computed by the VGG16 network, and similar visual attributes as computed by the ResNet-101.

From [Tables 5 and 6](#) it can be shown that our network perform favorably for each metric, both for RAP and AiC datasets. From [Fig. 6](#) it also emerges that our method, despite not using attention mechanisms, is able to detect and to remove the occlusion, with no external additional information. Furthermore, differently from the works by [Fabbri et al. \(2017\)](#) and [Isola et al. \(2017\)](#), our method learns to transfer with no alterations the portion of images that are not occluded.

## 6. Conclusions

In this work, we presented the use of GANs for image enhancing in people attributes classification. Our generator network has been designed to overcome a common problem in surveillance scenarios, namely people occlusion. Experiments have shown that jointly enhancing images before feeding them to an attribute classification network can improve the results even when input images are affected by this issue. We think that this line of work can foster research about the problem of attribute classification in surveillance contexts, where camera resolution and positioning cannot be neglected.

## Acknowledgments

The work is supported by the Italian MIUR, Ministry of Education, Universities and Research, under the project COSMOS PRIN 2015 programme 201548C5NT. We also gratefully acknowledge the support of Panasonic Silicon Valley Lab and Facebook AI Research with the donation of the GPUs used for this research.

## References

- Chen, C.Y., Grauman, K., 2014. Inferring unseen views of people. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 2003–2010.
- Coppi, D., Calderara, S., Cucchiara, R., 2016. Transductive people tracking in unconstrained surveillance. *IEEE Trans. Circuits Syst. Video Technol.* 26 (4), 762–775. <http://dx.doi.org/10.1109/TCSVT.2015.2416555>.
- Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L., 2009. ImageNet: A large-scale hierarchical image database. In: CVPR09.
- Deng, Y., Luo, P., Loy, C.C., Tang, X., 2014. Pedestrian attribute recognition at far distance. In: Proceedings of the 22nd ACM International Conference on Multimedia.
- Fabbri, M., Calderara, S., Cucchiara, R., 2017. Generative adversarial models for people attribute recognition in surveillance. In: *Advanced Video and Signal Based Surveillance. AVSS, IEEE International Conference on. IEEE*.
- Fabbri, M., Lanzi, F., Calderara, S., Palazzi, A., Vezzani, R., Cucchiara, R., 2018. Learning to detect and track visible and occluded body joints in a virtual world, arXiv preprint [arXiv:1803.08319](https://arxiv.org/abs/1803.08319).
- Ghodrati, A., Jia, X., Pedersoli, M., Tuytelaars, T., 2015. Towards automatic image editing: Learning to see another you, arXiv preprint [arXiv:1511.08446](https://arxiv.org/abs/1511.08446).
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y., 2014. Generative adversarial nets. In: *Advances in Neural Information Processing Systems 27*. Curran Associates, Inc., pp. 2672–2680.

- Guler, R.A., Neverova, N., Kokkinos, I., 2018. DensePose: Dense Human Pose Estimation in the Wild.
- He, K., Gkioxari, G., Dollár, P., Girshick, R., 2017. Mask r-cnn. In: *Computer Vision. ICCV, 2017 IEEE International Conference on. IEEE*, pp. 2980–2988.
- He, K., Zhang, X., Ren, S., Sun, J., 2016. Deep residual learning for image recognition. In: *CVPR. IEEE Computer Society*, pp. 770–778.
- Herath, S., Harandi, M., Porikli, F., 2017. Going deeper into action recognition: A survey. *Image Vis. Comput.* 60, 4–21, *Regularization Techniques for High-Dimensional Data Analysis*.
- Huang, R., Zhang, S., Li, T., He, R., et al., 2017. Beyond face rotation: Global and local perception gan for photorealistic and identity preserving frontal view synthesis, arXiv preprint [arXiv:1704.04086](https://arxiv.org/abs/1704.04086).
- Isola, P., Zhu, J., Zhou, T., Efros, A.A., 2017. Image-to-image translation with conditional adversarial networks. In: *2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21–26, 2017*, pp. 5967–5976, <http://dx.doi.org/10.1109/CVPR.2017.632>.
- Johnson, J., Alahi, A., Fei-Fei, L., 2016. Perceptual losses for real-time style transfer and super-resolution. In: *Computer Vision - ECCV 2016 - 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part II*, pp. 694–711.
- Kingma, D.P., Welling, M., 2013. Auto-encoding variational bayes, arXiv preprint [arXiv:1312.6114](https://arxiv.org/abs/1312.6114).
- Kupyn, O., Budzan, V., Mykhailych, M., Mishkin, D., Matas, J., 2017. DeblurGAN: Blind motion deblurring using conditional adversarial networks. *CoRR* abs/1711.07064.
- Lassner, C., Pons-Moll, G., Gehler, P.V., 2017. A generative model of people in clothing. In: *Proceedings of the IEEE International Conference on Computer Vision*, vol. 6.
- Ledig, C., Theis, L., Huszar, F., Caballero, J., Cunningham, A., Acosta, A., Aitken, A.P., Tejani, A., Totz, J., Wang, Z., Shi, W., 2017. Photo-Realistic Single Image Super-Resolution Using a Generative Adversarial Network. In: *2017 IEEE Conference on Computer Vision and Pattern Recognition. CVPR 2017, Honolulu, HI, USA, July 21–26, 2017*, pp. 105–114, <http://dx.doi.org/10.1109/CVPR.2017.19>.
- Li, D., Chen, X., Huang, K., 2015. Multi-attribute learning for pedestrian attribute recognition in surveillance scenarios. In: *2015 3rd IAPR Asian Conference on Pattern Recognition. ACPR*, pp. 111–115.
- Li, D., Zhang, Z., Chen, X., Ling, H., Huang, K., 2016. A richly annotated dataset for pedestrian attribute recognition, preprint [arXiv:1603.07054](https://arxiv.org/abs/1603.07054).
- Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L., 2014. Microsoft coco: Common objects in context. In: *European Conference on Computer Vision*. Springer, pp. 740–755.
- Liu, X., Zhao, H., Tian, M., Sheng, L., Shao, J., Yan, J., Wang, X., 2017. HydraPlus-Net: Attentive deep features for pedestrian analysis. In: *Proceedings of the IEEE International Conference on Computer Vision*, pp. 1–9.
- Ma, X., Zhu, X., Gong, S., Xie, X., Hu, J., Lam, K.M., Zhong, Y., 2017. Person re-identification by unsupervised video matching. *Pattern Recognit.* 65, 197–210.
- Mahendran, A., Vedaldi, A., 2015. Understanding deep image representations by inverting them, *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 5188–5196.
- Mirza, M., Osindero, S., 2014. Conditional generative adversarial nets, arXiv preprint [arXiv:1411.1784](https://arxiv.org/abs/1411.1784).
- Ouyang, W., Zeng, X., Wang, X., 2016. Partial occlusion handling in pedestrian detection with a deep model. *IEEE Trans. Circuits Syst. Video Technol.* 26 (11), 2123–2137.
- Pan, J., Hu, B., 2007. Robust occlusion handling in object tracking. In: *2007 IEEE Conference on Computer Vision and Pattern Recognition*, pp. 1–8, <http://dx.doi.org/10.1109/CVPR.2007.383453>.
- Pathak, D., Krahenbuhl, P., Donahue, J., Darrell, T., Efros, A.A., 2016. Context encoders: Feature learning by inpainting. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 2536–2544.
- Radford, A., Metz, L., Chintala, S., 2015. Unsupervised representation learning with deep convolutional generative adversarial networks, arXiv preprint [arXiv:1511.06434](https://arxiv.org/abs/1511.06434).

- Reed, S.E., Akata, Z., Mohan, S., Tenka, S., Schiele, B., Lee, H., 2016a. Learning what and where to draw. In: *Advances in Neural Information Processing Systems*. pp. 217–225.
- Reed, S., Akata, Z., Yan, X., Logeswaran, L., Schiele, B., Lee, H., 2016b. Generative adversarial text to image synthesis, *arXiv preprint arXiv:1605.05396*.
- Reed, S., van den Oord, A., Kalchbrenner, N., Bapst, V., Botvinick, M., de Freitas, N., 2016c. Generating Interpretable Images with Controllable Structure.
- Rezende, D.J., Mohamed, S., Wierstra, D., 2014. Stochastic backpropagation and approximate inference in deep generative models, *arXiv preprint arXiv:1401.4082*.
- Ronneberger, O., Fischer, P., Brox, T., 2015. U-Net: Convolutional networks for biomedical image segmentation. In: *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*. In: LNCS, vol. 351, Springer, pp. 234–241.
- Salimans, T., Goodfellow, I., Zaremba, W., Cheung, V., Radford, A., Chen, X., 2016. Improved techniques for training gans. In: *Advances in Neural Information Processing Systems*. pp. 2234–2242.
- Subramaniam, A., Chatterjee, M., Mittal, A., 2016. Deep neural networks with inexact matching for person re-identification. In: Lee, D.D., Sugiyama, M., Luxburg, U.V., Guyon, I., Garnett, R. (Eds.), *Advances in Neural Information Processing Systems* 29. Curran Associates, Inc., pp. 2667–2675.
- Sudowe, P., Spitzer, H., Leibe, B., 2015. Person attribute recognition with a jointly-trained holistic CNN model. In: *ICCV Workshops*. IEEE Computer Society, pp. 329–337.
- op het Veld, R.M., Wijnhoven, R., Bondarev, Y., et al., 2015. Detection and handling of occlusion in an object detection system. In: *Video Surveillance and Transportation Imaging Applications 2015*, vol. 9407, International Society for Optics and Photonics, p. 94070N.
- Wang, X., Hou, Z., Yu, W., Pu, L., Jin, Z., Qin, X., 2018a. Robust occlusion-aware part-based visual tracking with object scale adaptation. *Pattern Recognit.* 81, 456–470.
- Wang, C., Xu, C., Wang, C., Tao, D., 2018b. Perceptual adversarial networks for image-to-image transformation. *IEEE Trans. Image Process.* 27 (8), 4066–4079.
- Yan, X., Yang, J., Sohn, K., Lee, H., 2016. Attribute2image: Conditional image generation from visual attributes. In: *European Conference on Computer Vision*. Springer, pp. 776–791.
- Yang, C., Lu, X., Lin, Z., Shechtman, E., Wang, O., Li, H., 2017. High-resolution image inpainting using multi-scale neural patch synthesis. In: *The IEEE Conference on Computer Vision and Pattern Recognition*. CVPR, vol. 1, (2), p. 3.
- Yang, J., Reed, S.E., Yang, M.H., Lee, H., 2015. Weakly-supervised disentangling with recurrent transformations for 3D view synthesis. In: *Advances in Neural Information Processing Systems*. pp. 1099–1107.
- Yeh, R.A., Chen, C., Lim, T.Y., Schwing, A.G., Hasegawa-Johnson, M., Do, M.N., 2017. Semantic image inpainting with deep generative models. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 5485–5493.
- Yim, J., Jung, H., Yoo, B., Choi, C., Park, D., Kim, J., 2015. Rotating your face using multi-task deep neural network. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 676–684.
- Zhao, B., Wu, X., Cheng, Z.Q., Liu, H., Jie, Z., Feng, J., 2017. Multi-view image generation from a single-view, *arXiv preprint arXiv:1704.04886*.
- Zhuo, J., Chen, Z., Lai, J., Wang, G., 2018. Occluded person re-identification. *CoRR abs/1804.02792*.