

Warp and Learn: Novel Views Generation for Vehicles and Other Objects

Andrea Palazzi, Luca Bergamini, Simone Calderara, Rita Cucchiara

Abstract—In this work we introduce a new self-supervised, semi-parametric approach for synthesizing novel views of a vehicle starting from a single monocular image. Differently from *parametric* (i.e. entirely learning-based) methods, we show how *a-priori* geometric knowledge about the object and the 3D world can be successfully integrated into a deep learning based image generation framework. As this geometric component is not learnt, we call our approach *semi-parametric*. In particular, we exploit man-made object symmetry and piece-wise planarity to integrate rich *a-priori* visual information into the novel viewpoint synthesis process. An Image Completion Network (ICN) is then trained to generate a realistic image starting from this geometric guidance. This careful blend between parametric and non-parametric components allows us to i) operate in a real-world scenario, ii) preserve high-frequency visual information such as textures, iii) handle truly arbitrary 3D roto-translations of the input and iv) perform shape transfer to completely different 3D models. Eventually, we show that our approach can be easily complemented with synthetic data and extended to other rigid objects with completely different topology, even in presence of concave structures and holes (e.g. chairs). A comprehensive experimental analysis against state-of-the-art competitors shows the efficacy of our method both from a quantitative and a perceptive point of view.

Supplementary material, animated results, code and data are available at:
<https://github.com/ndrplz/semiparametric>

Index Terms—Vehicle novel view generation, Novel viewpoint synthesis, Texture transfer, Shape transfer, Semi-parametric.

1 INTRODUCTION

HOW would you see an object from another point of view? Given a single view of an object in the world, predicting how it would look like from arbitrarily different viewpoints is definitely non-trivial for both humans and machines. Still, people with a good *visual-spatial intelligence* [13] can easily imagine objects' rotation, zoom and translation shifts, especially if objects have a well known shape and feature some degree of symmetry. Indeed, humans have been shown to perform mental transformations for decision-taking about their surrounding environment [50], [2].

Machines are still far from this level of intelligence. Still, powerful *parametric* (i.e. entirely learning-based) deep learning models [27], [15] made it possible to frame the generation of novel viewpoints as a *conditioned image synthesis* problem. However, this is an holistic approach that under-exploits the fact that man-made objects 3D models are roughly distributed according to few prototypes (e.g. sedan, VAN, pick-up, truck etc. for vehicles). Up to now, the generation is constrained to be visually plausible with almost no geometric support from prototypes' shapes [58], [42]. Furthermore, even though generated images may look realistic *per se*, fine-grained visual appearance characterizing the particular object instance (e.g. texture) is often lost due to its high frequency which hardly survives being encoded through a deep network [58], [79], [10]. Eventually, most methods are supervised on the target image, which is seldom available for real data. Also, vast amount of data are required for the network to generalize to arbitrary transformations (i.e. a sufficient number of images for every possible viewpoint). This constrains many methods to be trained solely on synthetic data

and to be restricted to a discrete set of viewpoints: in fact, most recent works only handle a small set of transformations (e.g. fixed radius rotations) [69], [58], [76], [42].

At the same time, an independent line of research has shown that a *non-parametric* approach can be a viable path for photorealism, as also pointed out by Qi *et al.* [45]. For instance, new images can be generated by collaging [16], [28], [6], [21] or by leveraging multiple photographs to synthesize novel views via image-based rendering [4], [39], [18], [40]. Still, these methods require a large amount of data at test time: entire image banks for collaging, multiple photographs and depth data for image-based rendering.

In this work we propose a new approach - inherently *semi-parametric* - being based on both learning and geometry, self-supervised and efficient to be used in real-time. By taking the best from both worlds, we exploit geometric constraints to roughly sketch the target shape of the object and its textures while still relying on deep view synthesis to refine the generated view. The rationale behind this work is that many man-made objects adhere to a-priori geometric rules: vehicles in particular, exhibit a symmetric, piece-wise planar structure. Therefore, those properties may be exploited to approximately represent them by a small set of piece-wise planar patches, which can be warped almost exactly from source to destination viewpoint via a symmetry-aware homography transformation. Although these warped patches provide a rich hint about the visual content of the target viewpoint, they are far from being useful on their own. Thus, a fully-convolutional network is seeded with these patches along with a 2.5D CAD-rendered sketch to be used as guidance; it is then trained in a self-supervised manner to discriminate which part of the image must be completed or in-painted for the result to look realistic (see Fig. 3).

All authors are with University of Modena and Reggio Emilia, Italy.



Fig. 1. We propose a semi-parametric framework to generate realistic novel views of a vehicle and / or to transfer its appearance to different models.

We tested our solution in particular in vehicle generation due to their ubiquity in urban scene understanding applications [70], [59], [12], [49], [34], [30]. Moreover, to highlight that a decomposition in planar patches holds for different types of rigid objects, we evaluate our semi-parametric framework on both convex objects (*vehicles*) and concave ones (*chairs*). We leave as future work the analysis of a broader set of object categories.

In summary, our main contributions follow:

- We propose an original formulation of the problem of object novel viewpoint synthesis in a semi-parametric setting. Loose geometrical assumptions about the object shape provide rich hints about its appearance (*non-parametric*); this information guides a fully-convolutional network (*parametric*) in the generation process.
- We design our model to be trainable on existing datasets for 3D object detection in a *self-supervised* manner, with no need for paired source/target viewpoint images. Furthermore, we leverage 2D keypoints for real-world images where foreground segmentation is not provided.
- We demonstrate how our method excels in preserving visual details (e.g. texture) and in performing realistic shape transfer to completely different 3D models, while still being resilient to a much wider range of 3D transformations than competitors.

Our method can be employed to generate realistic novel views of an object from an arbitrary zoom, viewpoint and distance, as depicted in Figures 1, 4, 6, 7, 9, 11, 12. This could enable the generation of novel views from arbitrary virtual cameras in an urban scene, with applications in surveillance, vehicle re-identification and forensics. Also, our approach allows a disentangled editing of object shape and appearance (i.e. shape can be changed while preserving appearance or the other way around). This enables applications in interactive 3D manipulation and design, as well as data augmentation (Fig. 13).

A thorough experimental analysis is conducted comparing

our proposal with state-of-the-art methods, considering both the quantitative and the perceptual point of view.

2 RELATED WORK

View synthesis In just few years, the widespread adoption of deep generative models [27], [15] has led to astounding results in different areas of image synthesis [46], [1], [72], [24], [61], [71]. In this scenario, conditional GANs [35] have been demonstrated to be a powerful tool to tackle image-to-image translation problems [22], [77], [78], [8]. Hallucinating novel views of the subject of a photo can be naturally framed as an image-to-image translation problem. For human subjects, this has been cast to predicting the person’s appearance in different poses [33], [52], [73], [10]. Fashion and surveillance domains drew most of the attention, with much progress enabled by large real-world datasets providing multiple views of the same subject [31], [74].

For rigid objects instead, this task is usually referred to as *novel 3D view synthesis* and additional assumptions such as object symmetry are taken into account. In point of fact, symmetry is the most common assumption [76], [42], [79] to synthesise disoccluded portions of the object. Starting from a single image, Yang *et al.* [69] showed how a recurrent convolutional network can be trained via curriculum-learning to perform out-of-plane object rotation. In a similar setting Tatarchenko *et al.* [58] predicted both object appearance and depth map from different viewpoints. Successive works [76], [42] trained a network to learn a symmetry-aware appearance flow, re-casting the remaining synthesis as a task of image completion; [56] extends this approach to the case in which $N > 1$ input viewpoints are available. However, all these works [69], [58], [76], [42], [56] assume the target view to be known at training time. As this is not usually the case in the real-world, these approaches limit themselves to train solely on synthetic data and exhibit limited generalization in a real-world scenario. The recent work by Zhu *et al.* [79] exploits cycle consistency losses to overcome the need of paired data, thus training on datasets of segmented real-world cars and chairs they gathered for the purpose. Although that work shows more realistic results, it requires pixel-level segmentation




Fig. 2. We model a rigid object with a small of piece-wise planar patches, whose vertices are defined by 2D keypoints. We also include a small central crop as appearance prior to carry low-frequency information.

for each class of interest. In contrast, we show that already available datasets for object 3D pose estimation [67], [66] can be used for this purpose, despite the extremely rough alignment between the annotated model and the image.

3D Shape reconstruction Few recent works [63], [79] have shown that the use of 2.5D sketches can be a viable path to bridge the gap between synthetic and real-world data. In particular, in Zhu *et al.* [79] the 2.5D sketch consists of both a silhouette and a depth image rendered from a learnt low-resolution voxel grid by means of a differentiable ray-tracer. While this method is appealing for its geometrical guarantees, it is limited by a number of factors: i) it requires a custom *differentiable* ray-tracing module; ii) footprint of voxel-based representations scales with the cube of the resolution despite most of the information lying on the surface [54], [41]; iii) errors in the 3D voxel grid naturally propagate to the 2.5D sketch. We also follow this line of work to provide soft 3D priors to the synthesis process. However, in our semi-parametric setting 2.5D sketches are additional inputs which can be rendered from arbitrary viewpoints using standard rendering engines.

Nonparametric view synthesis In the interactive editing setting, recent works [25], [47] have shown astounding results by keeping the human in the loop and assuming a perfect (even part-level) alignment between the 3D model and the input image. However, as pixels are warped from the input to the target view [47] it is not feasible to perform shape transfer to a completely different model. Moreover, the time required to synthesise the output is still far from real-time (few seconds). On the opposite, our method enables disentangled shape and appearance transfer in real-time, with only a coarse alignment between the input image and the 3D model.

In the scenario of image synthesis from semantic layout the recent work of Qi *et al.* [45] has shown that non-parametric components (i.e. a memory bank of image segments) can be integrated in a parametric image synthesis pipeline to produce impressive photo-realistic results. Despite our different setting, we similarly rely on image patches to provide hints to the Image Completion Network (ICN); however, our patches are not queried from a database but warped directly from the input view.

After our initial submission, several new related works have been published by the scientific community. We include here a non exhaustive list of them; [36], [68], [55], [37], [38], [7].

3 MODEL

Our model generates novel views of objects in a semi-parametric setting - relying on both geometry and learning. To this end, 2D keypoints and additional 3D information are extracted from a single view of the object. Keypoints are used as a proxy to describe 2D geometrical abstractions of the 3D shape (i.e. planar patches), which are transformed to the novel viewpoint. Eventually, a convolutional neural network seamlessly fuses this prior information to generate a realistic image from the novel view.

More into details, our model takes as input an image depicting a single object $\mathbf{x}_s \in \mathbb{R}^{H \times W \times 3}$ viewed from the source viewpoint $\mathbf{V}_s \in \mathbb{R}^{4 \times 4}$, its 2D keypoints \mathcal{K}_s and its associated 3D CAD model $C \in \mathbb{R}^{faces \times 3 \times 3}$ having 3D keypoints \mathcal{K}_{3D} .

Training (Fig. 3, top) is performed in a *self-supervised* fashion maximizing the consistency between the input image and the generated one when projected onto the source viewpoint, with no need for coupled images from the two viewpoints as supervision. Given \mathbf{x}_s and \mathcal{K}_s , planar patches are extracted (Sec. 3.1). The patches are then projected to the target viewpoint \mathbf{V}_d through an intermediate view (Sec. 3.2) according to a visibility model (Sec. 3.3). The 3D model C is also rendered from the target viewpoint \mathbf{V}_d to get a 2.5D sketch of the object (Sec. 3.4). Eventually, the image completion network (ICN) starts from these visual seeds to generate a realistic final image (Sec. 3.5, 3.6).

Inference (Fig. 3, bottom) follows a similar flow. However, in this case only 2D keypoints \mathcal{K}_s and source viewpoint \mathbf{V}_s are needed; the 3D model can be either inferred from the input or arbitrarily selected to perform shape transfer.

Without loss of generality, in this work we rely on ground truth data whenever possible, as our focus lies on the overall viewpoint generation pipeline. Off-the-shelf detectors [17], [44], [60] can be used to provide these information in an in-the-wild scenario.

3.1 Keypoint-based decomposition into planar patches

We leverage 2D keypoints to approximate the visible shape of the object with a simple polyhedron with a small set of faces, as depicted in Fig. 2. Since keypoints mark characteristic locations in the object shape (e.g. corners), a face defined from at least three of those could carry a perceptual / semantic meaning (e.g. the roof of a car). Exploiting 2D keypoints to find object faces is appealing for a number of reasons. First, this makes straightforward computing the homography matrix between planes in different viewpoints (see Sec. 3.2). Furthermore, a number of datasets provide object landmark annotations in real-world scenarios (e.g. [29], [67], [66], [62], [64], [20]) and solid keypoints detection methods exist [60], [44], [17].

Specifically, for each source image $\mathbf{x}_s \in \mathbb{R}^{H \times W \times 3}$ an array of 2D keypoints $\mathcal{K}_s \in \mathbb{R}^{|\mathcal{K}_s| \times 2}$ is available, being $|\mathcal{K}_s|$ the category-specific number of keypoints (we set $|\mathcal{K}_s| = 12$ for vehicles). From these a set of planar patches are extracted according to the geometry of the object:

$$\mathcal{P}_s = \{\mathbf{p}_s^{(0)}, \mathbf{p}_s^{(1)}, \dots, \mathbf{p}_s^{(|\mathcal{P}|)}\}, \quad \mathbf{p}_s^{(i)} = \Gamma(\mathbf{k}_s^{(i)}) \quad (1)$$

where each patch $\mathbf{p}_s^{(i)}$ is defined as the convex hull Γ of the subset of keypoints $\mathbf{k}_s^{(i)} \subseteq \mathcal{K}_s$. Prior knowledge about the object class can be leveraged to choose the planar patches \mathcal{P}_s : for instance, we choose roof, left, right, front and back sides for vehicles.

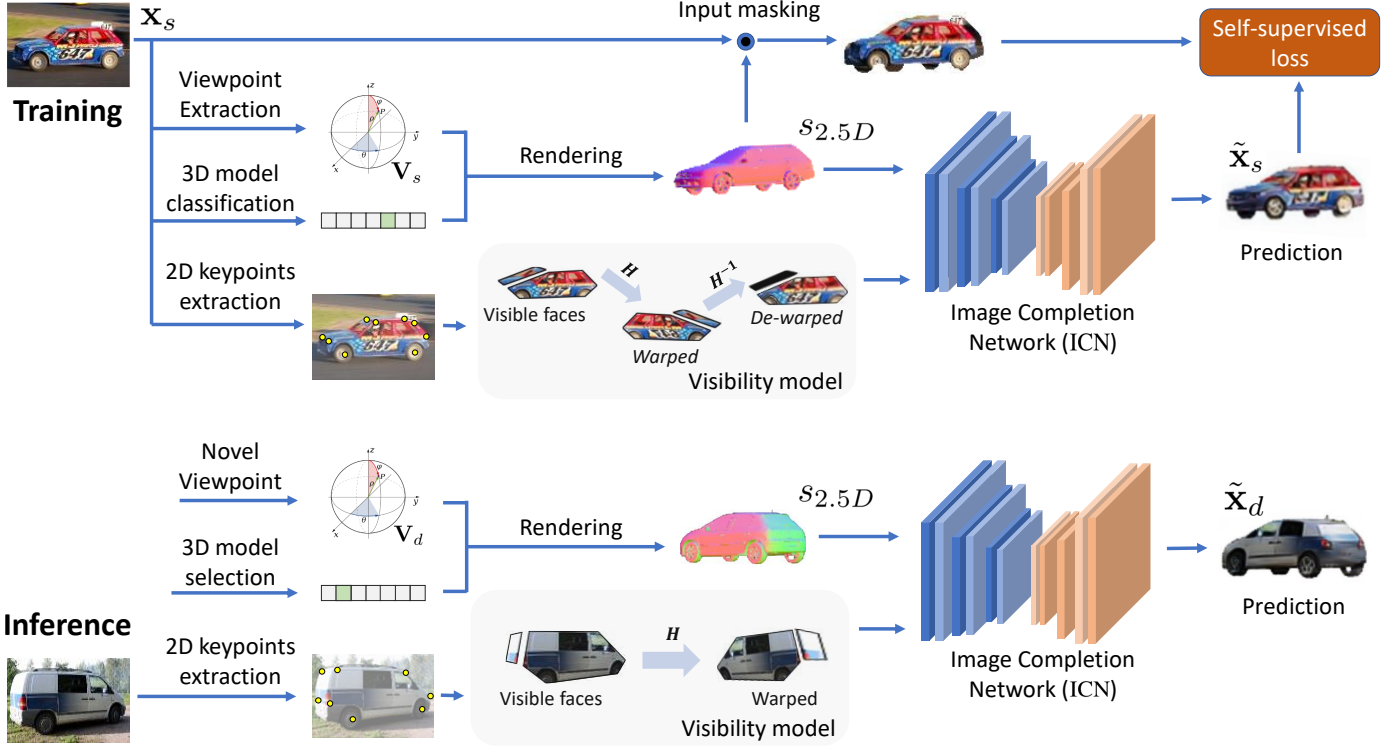


Fig. 3. Model architecture overview. Approximately planar patches are extracted from the 2D keypoints locations. The Image Completion Network (ICN) uses the synthetic 2.5D sketches as templates to reconstruct object's appearance from the patches in a self-supervised fashion. During training, input patches are warped forth and back to a randomly sampled viewpoint to enforce resilience against homography issues that are likely to be encountered at test time. During inference, novel views of the input object are synthesised by providing the ICN a novel viewpoint and a (possibly different) rendered 3D model to be used as shape guideline.

3.2 Warping and dewarping

Warping patches Source patches \mathcal{P}_s are warped to the destination viewpoint to get a set of warped patches \mathcal{P}_d that are employed to bootstrap the novel viewpoint synthesis. To this end, we define the destination viewpoint $\mathbf{V}_d \in \mathbb{R}^{4 \times 4}$ to be an arbitrary rigid transformation of the camera:

$$\mathbf{V}_d = \begin{bmatrix} \mathbf{R} & \mathbf{t} \\ \mathbf{0}^T & 1 \end{bmatrix} \quad (2)$$

Locations of 2D keypoints $\mathcal{K}_d \in \mathbb{R}^{|\mathcal{K}_s| \times 2}$ in the novel viewpoint can be now computed by using the classical pinhole camera model as:

$$\mathbf{k}_d^{(i)} = \begin{bmatrix} f & 0 & c_x \\ 0 & f & c_y \\ 0 & 0 & 1 \end{bmatrix} \mathbf{V}_d^{-1} \mathbf{k}_{3D}^{(i)} \quad (3)$$

where $\mathbf{k}_{3D}^{(i)}$ is the i^{th} 3D keypoint in the CAD model and c_x, c_y are the principal point coordinates. As the focal f is unknown, we set it to an high value to minimize perspective effects; we choose $f = 3000$ as in Pascal3D+ [67]. A set of homography transformations \mathbf{H} relating planar surfaces in the two views can be estimated from correspondences between \mathcal{K}_s and \mathcal{K}_d . In this way patches in the destination viewpoint (*warped patches* from now on) can be computed as:

$$\mathcal{P}_d = \{\mathbf{p}_d^{(0)}, \mathbf{p}_d^{(1)}, \dots, \mathbf{p}_d^{(|\mathcal{P}|)}\}, \quad \mathbf{p}_d^{(i)} = \mathbf{H}^{(i)}(\mathbf{p}_s^{(i)}) \quad (4)$$

De-warping patches Since real-world datasets do not provide paired views, it is not possible to supervise the destination image

$\tilde{\mathbf{x}}_d$; hence, we propose to train the ICN network in a *self-supervised* manner. A straightforward approach would be to reconstruct \mathbf{x}_s from \mathcal{P}_s . Nonetheless, this would create a distribution shift between the data fed to the network during training and inference stages. In fact, while \mathcal{P}_s is perfectly aligned with \mathbf{x}_s , \mathcal{P}_d might be affected by homography failures and interpolation errors among other issues. For example, when the destination area is smaller than the source one many source pixel land onto the same destination pixel, and the inverse warping cannot recover all the information in the original patch. To alleviate this shift, we train the network to reconstruct the image \mathbf{x}_s from a third set of patches (called *dewarped patches* in what follows):

$$\tilde{\mathcal{P}}_s = \{\tilde{\mathbf{p}}_s^{(0)}, \tilde{\mathbf{p}}_s^{(1)}, \dots, \tilde{\mathbf{p}}_s^{(|\mathcal{P}|)}\}, \quad \tilde{\mathbf{p}}_s^{(i)} = (\mathbf{H}^{(i)})^{-1}(\mathbf{p}_d^{(i)}) \quad (5)$$

During training, patches are warped towards a random viewpoint sampled from the training set distribution before being warped back to the source one. In this way the network learns to cope with possible transformation errors and cannot simply short-circuit input patches to the output. The importance of this *dewarping trick* for a well-behaved network training is highlighted in Sec. 4.1.

3.3 Visibility model

Whenever a 3D object is projected into a 2D image, self-occlusions almost inevitably arise. Consequently, not all planar patches into the set \mathcal{P}_s are effectively visible. Were they to be warped regardless of their visibility, following parts of our architecture would require to discern which of them to keep or discard. Furthermore, when warping between \mathcal{P}_s and \mathcal{P}_d the visibility of some of those planes may vary. To take these

dynamics into account, we first render the object 3D model from the camera viewpoint to obtain the 3D planes corresponding to the detected 2D patches. The z-buffer computed through ray-casting is then exploited to filter the patches which are not visible from the source viewpoint. These are “dropped”, in the sense that they are zeroed before feeding them to the ICN. As during training the intermediate viewpoint \mathbf{V}_d is sampled randomly, the warping-deforming phase results in a random dropout at patch-level, where the chance of drop is inversely proportional to the frequency of visibility of the patch. This forces the network to hallucinate missing patches during training, thus improving generalization when source and destination viewpoint differ.

3.4 Leveraging 2.5D sketches

While image patches carry rich information about the appearance of the object, they bear few cues about the object shape. In other words, visual aspect and shape are disentangled by design. This is a desirable property enabling multiple applications which require to change one of the two while keeping the other fixed. In this section we propose a method to constrain the synthesised object shape. Let

$$\mathcal{C} = \{C^{(0)}, C^{(1)}, \dots, C^{(|\mathcal{C}|)}\}, \quad C^{(i)} \in \mathbb{R}^{f_{\text{faces}} \times 3 \times 3} \quad (6)$$

be the set of 3D CAD models which approximate the intra-class variation for the current object class, each $C^{(i)}$ being a 3D mesh composed of f faces. The number of CADs $|\mathcal{C}|$ needed to cover the intra-class variation reasonably depends on the object category, but it is often relatively low (e.g. $|\mathcal{C}| = 10$ for the vehicle class in the Pascal3D+ dataset [67]). Each training example i is thus composed by an image $\mathbf{x}^{(i)}$ and its associated viewpoint $\mathbf{V}^{(i)}$ and CAD index $\alpha \in \{0, 1, \dots, |\mathcal{C}|\}$ which can be possibly selected through a classifier. Therefore, a virtual camera can be used to render the CAD $C^{(\alpha)}$ from viewpoint $\mathbf{V}^{(i)}$. In particular, following [63], we render the 2.5D sketch of CAD surface normals:

$$s_{2.5D}(C^{(\alpha)}, \mathbf{V}^{(i)}) \in \mathbb{R}^{H \times W \times 3} \quad (7)$$

which provides rich information about the object’s 3D shape. During training, this 2.5D sketch is fed to the ICN together with de-warped patches $\tilde{\mathcal{P}}_s$ to reconstruct \mathbf{x}_s .

3.5 Appearance prior

Our method relies on warped patches to transfer the object appearance from a source to a destination viewpoint. Still it might happen that viewpoints \mathbf{V}_s and \mathbf{V}_d are so far apart (e.g. front to back) that an object shares no visible faces across the two even with symmetry constraints. To alleviate this issue, we crop from the center of the input image \mathbf{x}_s a small patch c_s with side 10% of the image size and give it as an additional input to the image completion network as a prior knowledge about the rough object appearance in absence of other hints (depicted in Fig. 2). The network can extract from this crop coarse information about the object visual aspect (e.g. the average color) as a prior to cope with large changes between viewpoints.

3.6 Image Completion Network

The Image Completion Network (ICN) $g(\cdot \mid \theta)$ is a fully convolutional network parametrized by θ trained to reconstruct a realistic

image \mathbf{x}_s from de-warped patches $\tilde{\mathcal{P}}_s$, 2.5D sketch $s_{2.5D}$ and appearance prior c_s :

$$\tilde{\mathbf{x}}_s = g(\tilde{\mathcal{P}}_s, s_{2.5D}(C^{(\alpha)}, \mathbf{V}_s), c_s \mid \theta) \quad (8)$$

Architecture Our ICN features an encoder / decoder structure as in [77], [79]. The encoder is composed of 3 convolutional blocks to reduce the spatial resolution and 3 additional residual blocks applied to the lowest resolution code. Except for the first one, every convolution has kernel size 4x4 and it’s preceded by reflection padding, while ReLU activation and Instance normalization are applied afterwards. The decoder follows the same structure in reverse order. For the discriminator network we rely on the two-scale PatchGAN classifier [22], [77], [79].

Objective A number of recent works [23], [5], [45] indicate that loss functions based on high-level features extracted from pre-trained networks can lead to much more realistic results compared to naive *per-pixel* losses between the output and ground-truth image. Given a set of layers $\{\Phi_l\}$ from a network Φ and a training pair consisting of a real and a generated images $(\mathbf{x}_s, \tilde{\mathbf{x}}_s)$, we define the perceptual loss function as

$$\mathcal{L}_{\mathbf{x}_s, \tilde{\mathbf{x}}_s}^{VGG}(\theta) = \sum_l \lambda_l \|\Phi_l(\mathbf{x}_s) - \Phi_l(\tilde{\mathbf{x}}_s)\|_1. \quad (9)$$

We employ each second convolutional layer of each block in VGG-19 [53] as feature extractor Φ_l ; $\{\lambda_l\}$ is set as in [5].

As mentioned above, images generated from novel viewpoints $\tilde{\mathbf{x}}_d$ cannot be directly supervised if the dataset does not provide paired views. Nevertheless, we can still enforce the realism of ICN output in an adversarial fashion. Given a generic image $\tilde{\mathbf{x}}$ synthesised by ICN either in the source ($\tilde{\mathbf{x}}_s$) or the destination ($\tilde{\mathbf{x}}_d$) viewpoint, we set up a min-max game as follows:

$$\mathcal{L}_{\mathbf{x}_s, \tilde{\mathbf{x}}}^{adv} = \mathbb{E}_{\mathbf{x}_s} [\log D(\mathbf{x}_s)] + \mathbb{E}_{\tilde{\mathbf{x}}} [\log(1 - D(\tilde{\mathbf{x}}))] \quad (10)$$

where D is the discriminator network aiming to distinguish between real and synthesised images. Our total loss is defined as:

$$\mathcal{L} = \mathcal{L}_{\mathbf{x}_s, \tilde{\mathbf{x}}_s}^{VGG} + \gamma \mathcal{L}_{\mathbf{x}_s, \tilde{\mathbf{x}}}^{adv} \quad (11)$$

where γ modulates the contribution of the adversarial term.

4 EXPERIMENTS

The experiments we propose concern with the evaluation of the quality of the generated object images across different viewpoints. First, both visual comparison (Sec. 4.1) and quantitative experiments (Sec. 4.2) against state-of-the-art competitors are reported. Then, we keep the human in the loop by relying on human judgement to measure the output quality via A/B preference tests (Sec. 4.3). Eventually, we extend the evaluation to other classes (Sec. 4.4) and we investigate the contribution of complementary synthetic data for modelling extreme viewpoint changes (Sec. 4.5).

Datasets Although large-scale 3D shape repositories providing object geometries such as Princeton Shape Benchmark [51] and Shapenet [3] exist, they do not come with real-world images aligned. As we want to work with real-world data, we rely on Pascal3D+ [67], an in-the-wild 3D object detection dataset which augments the 12 rigid categories of the PASCAL Visual Object Classes (VOC) [11] with 3D annotation -roughly [57]- aligned. In particular, we use the *car* and *chair* subsets, which consist of around 5000 and 1500 images respectively.

Competitors We evaluate our method against six state-of-the-art works in the task of novel viewpoint synthesis. The first one is




Fig. 4. Results of 360° rotation. Our output is consistent for the whole rotation circle. Best viewed zoomed on screen.

	0°	30°	60°	90°	120°	150°	180°	210°	240°	270°	300°	330°	Avg
ours_{real+synth}	174.2	57.0	47.9	53.8	52.9	59.9	168.8	61.2	52.1	54.2	46.8	60.2	74.1
ours_{real}	178.9	54.6	51.0	53.2	56.8	58.1	202.5	62.2	54.1	49.8	46.3	56.1	77.0
MV2NV [56]	144.5	170.9	136.5	192.3	144.7	163.1	152.2	179.1	137.5	186.0	136.0	161.3	158.7
MV3D [58]	257.5	263.9	259.7	284.2	273.0	271.9	267.7	263.7	261.5	277.9	266.5	265.3	267.7
TVSN [42]	70.8	71.3	74.1	78.9	79.0	82.2	89.4	81.1	79.0	78.9	78.0	72.9	78.0
VUnet [10]	202.4	90.7	79.9	88.3	78.8	96.3	203.3	94.0	77.5	85.2	82.0	92.6	105.9
VON [79]	134.1	107.2	150.0	126.5	124.4	114.4	151.7	113.8	127.2	128.9	132.0	107.3	126.4
VON _{FT} [79]	165.2	100.6	137.8	125.9	137.6	108.1	190.5	155.7	134.8	123.1	117.1	102.1	133.2

TABLE 1

Fréchet Inception Distances [19] results for *car*. Each row lists the average distance between real and generated images for each method on the left. Results are reported from 12 evenly spaced azimuthal angles while rotating around the object at fixed elevation and radius. Details in Sec. 4.2.



Fig. 5. Comparison with ablated versions of the proposed method on Pascal3D+ test set. Better viewed zoomed on screen. Please refer to Sec. 4.1 for details.

VON [79], an adversarial learning framework in which object shape, viewpoint and texture are treated as three conditionally independent factors that contribute to the synthesis of the novel viewpoint. Since VON was originally trained on a custom car dataset collected by the authors [79], for a fair comparison we implement a second baseline VON_{FT} by fine-tuning their network on Pascal3D+. For both competitors we provide ground truth voxelized shapes from [79] matching the images' CADs, relying on the texture encoder network from [79] for extracting the textures from Pascal3D+ images. Third, we compare to VUnet [10], a state-of-the-art framework for conditional image generation based on variational autoencoder [27], which shows a good generalization capability across a variety of poses and viewpoints. In the authors' implementation [10] a U-net [48] architecture is fed with keypoint-based skeletons to perform pose-guided human generation. We re-train their model on Pascal3D+ to perform pose-guided object

generation. For this process, we feed their shape-encoding network with our 2.5D sketches rendered from the object CAD - which is a more informative signal than the one used in the original implementation (i.e. skeleton or edges) as it also includes the direction of the normal surfaces and fine-grained details from the 3D model.

We also compare with three recent *pairwise-trained* models: MV3D [58], TVSN [42] and MV2NV [56]. Pairwise methods share the need for both source and target pairs during training; thus we cannot re-train or finetune them on Pascal3D+ and we rely on pre-trained models released by the authors.

To maximize evaluation fairness, in what follows we only sample novel viewpoints rotating around the z-axis at fixed distance and elevation, which is the only setting handled by competitors. Still, our method can handle general roto-translations as well as variation in camera intrinsic. Fig. 7 and Fig. 9 show visual results for large viewpoint changes in both elevation and azimuth; even more extreme transformations are depicted in Fig. 12.

Implementation details The 2D bounding box of each example of Pascal3D+ [67] is padded to a squared aspect ratio and resized to 128x128 pixels. We work in LAB space relying on the training procedure from [65]. Following [60], [44] truncated and occluded objects are discarded, resulting in 4081 training and 1042 testing examples respectively. For vehicles, the ours_{real+synth} model is trained for 20 epochs with batch size 8. During training images undergo small random rotations, translations and shearing for data augmentation purposes. We use Adam [26] optimizer with constant learning rate $2e^{-4}$. Loss balancing term γ is set to 8. The code is developed in PyTorch [43]: we depend on Open3D library [75] for 3D data manipulation and rendering. Random search has been employed for hyper-parameter tuning. Without any optimization, inference for a 128x128 image takes $\sim 3ms$ on a NVIDIA GTX 1080, making it suitable for real-time applications.



Fig. 6. Visual results comparison with competitors on Pascal3D+ test set. Better viewed zoomed on screen. Please refer to Sec. 4.1 for details.




Fig. 7. Predictions of our model from different viewpoints. The geometry-aware design of our semi-parametric method allows the model to be resilient to large viewpoint variations, including rotation, elevation and camera distance.

4.1 Visual results

Our model produces high-quality results for a variety of camera viewpoints, preserving fine-grained object appearance, as it can be appreciated in Figures 1, 4, 5, 6, 7, 11 and 12.

Competitors The key differences between the proposed method and competitors can be appreciated in Fig. 6. From left to right. MV2NV [56] seems to suffer the most the reality gap as well as the lack of multiple views, leading to the worse quality results in our setting. MV3D [58] predicts at least a reasonable overall shape, although images are blurry to the point that in some cases one can barely recognize the vehicle. Results from TVSN [42] show the highest variability: while looking generally fine, they are disastrous for less common poses such as (b) and (e). The output from Visual Object Networks [79] (VON) is generally realistic, but hardly reflects the visual appearance of the input. Furthermore, both VON and VON_{FT} generator networks do not generalize to poses which are less common in the training set such as the frontal pose in (d). VUnet [10] suffers from blurred results typical from variational autoencoders [14]; also, due to skip connections, input appearance may leak to the output when the two viewpoints are very different (b). More generally, the drawbacks of a solely

learning-based viewpoint synthesis are evident in (a, c): complex textures cannot be recovered once compressed in a feature vector. **Ablation** Ablated versions of our model are shown in Fig. 5. The effect of removing the appearance prior is showcased in *No-prior* column. Without prior information, the ICN fails to infer the object appearance when no planar patch is provided, as shown in (b, e). Removing the adversarial term (*No-adv* column) results in slightly blurred outputs. We also investigate the aid of the *dewarping trick* presented in Sec. 3.2. In *No-dewarp* column, the ICN was trained to reconstruct the image from \mathcal{P}_s instead of $\tilde{\mathcal{P}}_s$ (see Sec. 3.1). As expected, despite the very low reconstruction error at training time (due to the similarity between \mathbf{x}_s and \mathcal{P}_s), the model fails to generalize to the synthesis of novel viewpoints where the textures \mathcal{P}_d are the result of an homography transformation. This highlights the importance of the *dewarping trick* for a well-behaved training. Eventually, *Sil-only* shows the ablated versions in which the input sketch is constituted only by 2D silhouette. Although results do not differ dramatically, it can be appreciated how the network benefits from additional information to resolve ambiguous situations such as self-occlusions (e) and details such as side windows, lights, wheels.

	Car (plain)	Car (textured)	Avg
ours > <i>VUnet</i> [10]	76.0%	85.0%	78.0%
ours > <i>VON</i> [79]	88.0%	98.0%	91.0%
ours > <i>VON_{FT}</i> [79]	96.0%	99.0%	97.0%

TABLE 2

Blind randomized A/B test results. Each row lists the percentage of workers who preferred the novel viewpoint generated with our method with respect to each baseline (chance is at 50%).

Shape transfer Visual results for shape transfer are showed in Fig. 1. In this setting, the network is requested to complete the warped faces \mathcal{P}_d using the 2.5D sketch rendered from a totally different CAD. It can be appreciated how novel viewpoints are still realistic, since the network exploits the 2.5D sketch to complete the warped appearance in a CAD-agnostic manner.

4.2 Metrics evaluation

Fréchet Inception Distance To quantitatively measure the similarity between generated and real images we rely on Fréchet Inception Distance (FID), which was shown to consistently correlate with human judgment [19], [32]. We employ activations from the last convolutional layer of an InceptionV3 model pretrained on ImageNet [9] as features. Assuming a multidimensional Gaussian distribution for these features, we compute the FID as follow:

$$\text{FID} = \|\mathbf{m} - \mathbf{m}_w\|_2^2 + \text{Tr}(\mathbf{C} + \mathbf{C}_w - 2(\mathbf{C}\mathbf{C}_w)^{1/2}) \quad (12)$$

Where \mathbf{m} , \mathbf{C} are the mean and covariance of the features extracted from Pascal3D+ data, while \mathbf{m}_w and \mathbf{C}_w are the corresponding statistics extracted from the generated images.

To enable the comparison with other works, we sample novel viewpoints while rotating around the object at fixed distance and elevation: results are reported in Table 1, binned in 12 equidistant azimuthal angles. Fréchet Inception Distance rewards the realism we can get with our semi-parametric approach; our method outperforms all competitors for the vast majority of viewpoints. In particular, our method preserves both low (i.e. shape) and high-frequency (i.e. texture) image statistics which are both contribute to the overall FID score.

4.3 Perceptual experiments

To assess the quality of our results also from a perceptual point of view, randomized A/B preference tests were performed by 43 human workers, following the experimental protocol of previous works [5], [45], [79]. Due to time constraint, for this phase we




Fig. 8. Results of time-limited A/B preference test against real images. Both *VON* and our method are resilient to human judgement over time. Green line denotes random chance. Please refer to Sec.4.3 for details.

select only the three most recent competitors, namely *VON* [79], *VON_{FT}* [79] and *VUnet* [10]. As we want to evaluate both the realism and the appearance coherence of our method, we perform two different tests.

View transfer coherence In the first setting, the subject is presented with three images: while the first one comes from Pascal3D+ test set, *A* and *B* depict a novel viewpoint of the object generated with two different methods. The human worker is then asked whether rotating the input object would better lead to *A* or *B*. Results reported in Tab. 2 indicate that our method is largely preferred to competitors, likely because of the built-in realism that comes from warping the original image. As a further analysis we split by manual annotation Pascal3D+ images into *plain* and *textured* sets, the latter set containing vehicles which feature characteristic textures. Table 2 highlights that workers expressed almost unanimous preference for our method on the *textured* set. The fact that human attention was caught by these appearance details highlights the importance of preserving fine-grained details in the synthesized output.

Output realism The second experiment consists of a two-alternative forced choice aimed at evaluating the relative realism of each method. Here the subject is presented with only two images for a determined amount of time. The worker is then asked which of the two appeared more realistic and the experiment is repeated by varying the amount of time. Results depicted in Fig. 8 follow two trends. On the one hand, workers clearly discern *VUnet* and *VON_{FT}* images from real ones as more time is available. *VUnet* is hurt by excessive blur and visual artifacts; *VON_{FT}* suffers from a severe loss of realism w.r.t. the original *VON* method, which may be related to the great variety of viewpoints in the Pascal3D+ dataset compared with the one used in Zhu *et al.* [79]. On the other hand, both *VON* and our method produce realistic images workers struggle to distinguish from the real ones even in 8000ms.

4.4 Evaluation on other classes

Similarly to [56], [42], [79], we test our method also on the *chair* subset from Pascal3D+ dataset, consisting of 1195 images annotated with 10 different CAD models to assess the generalisation capability of our method. As for vehicles, also for chairs we define a set of planes from 3D annotated keypoints to approximate the surface of the objects, namely: left, right, seat and back. Even though chairs often feature holes and slits leaking some of the background, we treat them as filled planes. This is due to having access to the foreground segmentation mask through the rendered CAD model, which can be used to mask out background areas. However, the very coarse alignment of the chairs CAD models in Pascal3D+ leads to an additional difficulty when training the ICN. Table 3 reports results against competitors in terms of Fréchet Inception Distance score: MV3D [58] is omitted since it only releases pre-trained model for cars. Our method outperforms all other competitors also for this class of objects; visual examples are reported in Fig. 11. While MV2NV [56] and TVSN [42] often struggle to generate the object from the correct viewpoint, *VON* [79] fails to transfer visual details in the final output. Although the generated image looks realistic, it doesn't resemble the input one. Contrarily, our method successfully generate realistic views of the input object, even under severe viewpoint transformations 12. Still, when background leaks in the input mask due to the CAD misalignment (first and fourth row) the final output quality decreases.

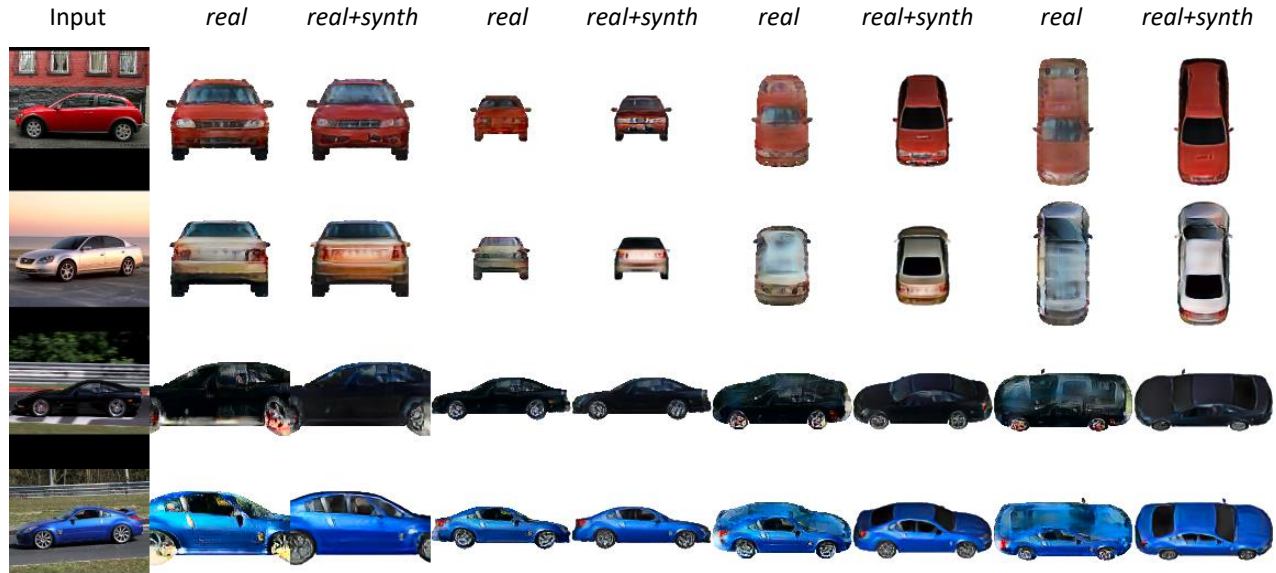


Fig. 9. Visual comparison showing the effect of adding synthetic data to Pascal3D+ training set. The ICN network trained on the mixture of the two domains performs significantly better under extreme viewpoint transformations. Please see supplementary material for more examples.

	0°	30°	60°	90°	120°	150°	180°	210°	240°	270°	300°	330°	Avg
ours_{real+synth}	83.9	57.9	66.5	120.1	82.5	76.1	102.3	77.2	83.4	117.0	65.2	60.9	82.8
ours _{real}	92.6	62.4	68.0	122.5	84.8	84.7	117.3	82.3	85.0	127.5	70.1	64.0	88.4
TVSN [42]	84.7	86.6	90.8	95.1	94.0	97.2	93.9	94.6	95.0	93.0	87.8	82.7	91.6
VON [79]	125.5	107.7	100.1	121.1	122.6	136.8	203.9	141.4	123.8	114.2	102.1	96.1	124.6
VUnet [10]	118.9	97.1	123.1	171.0	160.5	137.9	151.2	141.9	155.8	154.9	120.1	95.8	135.7
MV2NV [56]	180.3	168.0	178.3	187.8	184.1	177.8	184.7	184.3	200.2	191.0	184.7	166.4	182.3

TABLE 3

Fréchet Inception Distances [19] results for *chair* class. Each row lists the average distance between real and generated images for each method on the left. Results are reported from 12 evenly spaced azimuthal angles while rotating around the object at fixed elevation and radius.




Fig. 10. Viewpoints' distributions for real (a) and synthetic data (b) of the car class in Pascal3D+. Radii have been normalised to unit length for clearness. In red viewpoints with elevation lesser or equal than $\frac{\pi}{8}$ rad.

4.5 On the use of synthetic data

Even though the semi-parametric nature of our proposed method copes with a variety of viewpoint, it still rely on data to learn how to stitch together the warped patches. Therefore, for dramatic changes of viewpoint that are completely uncovered in the dataset, performances may drop significantly. As shown in Fig. 10 (a), Pascal3D+ viewpoints' distribution for the car set is profoundly skewed towards low elevation values and is polarized with regard to the azimuth to frontal and lateral views, reflecting how the images were acquired. As it is of great interest to produce realistic images from more varied viewpoints (e.g. bird's eye view), we include synthetic data in the training set for balancing

the viewpoints' distribution. To this end, we sample 59 models from the car synset in the ShapeNet [3] dataset and we annotate them with 3D keypoints to define the planes of interest. We then render 6950 images sampling viewpoints uniformly in a 3D semi-spheres around the origin. As shown in Fig. 10 (b) this viewpoints' distribution is much more uniformly distributed in terms of azimuth and elevation. We name ours_{real+synth} the ICN network trained on a mixture of synthetic and real data. We also experimented pre-training our model on synthetic data and fine-tuning on the real ones, although in our experience that policy led to worse results.

Visual comparison between ours_{real} and ours_{real+synth} is shown in Fig. 9. It can be seen how training only on Pascal3D+ entails artefacts for out-of-distribution viewpoints (e.g. bird's eye view). Conversely, combining the two domains the network learns from synthetic data a prior about the overall structure and color. We perform an analogous augmentation of the chair class. In this case we included 1858 freshly rendered synthetic images from 73 annotated models. As the rendered images are perfectly aligned with the foreground mask, this also contributes to reduce the planes' misalignment introduced by images from Pascal3D+.

5 CONCLUSIONS

In this work we introduced a novel formulation of the problem of object novel viewpoint synthesis in a semi-parametric setting. Notably, our model is designed to be trainable on existing datasets



Fig. 11. Visual results comparison with competitors for *chair* class on Pascal3D+ test set. Better viewed zoomed on screen.

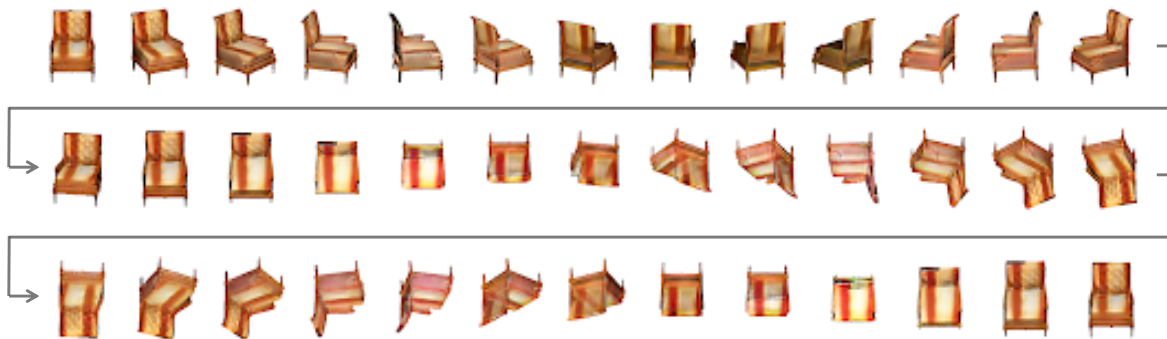


Fig. 12. Geometric guidance built in our semi-parametric model allows to perform extreme object transformations which are currently unfeasible for any fully-parametric method. In this case we are able to rotate the armchair upside-down, although this configuration never appears in the training set. Best viewed zoomed on screen.



Fig. 13. Artificial data created stitching generated vehicles onto Pascal3D+ [67] backgrounds.

for 3D object detection in a self-supervised manner, without the need for paired source/target viewpoint images - although it can be complemented with synthetic data. Non-parametric visual hints act as prior information to guide a deep parametric model for generating realistic images, disentangling by design appearance

and shape. This enables truly continuous manipulation of the view-point and shape transfer to different 3D models. As completing the image is much easier than generating it from scratch, we can train our ICN on just few thousands images from the Pascal3D+ dataset and still be able to generalize to unseen viewpoints. Although vehicles were the main focus of our work, we show that our framework is generic enough to handle rigid objects of completely different geometric structure such as chairs. Perceptual experiments results as well as image-quality metrics reward our method for its realism and the visual consistency of the synthesised object across arbitrary points of view.

ACKNOWLEDGEMENTS

This research was supported by MIUR PRIN project "*PREVUE: Prediction of activities and Events by Vision in an Urban Environment*", grant ID E94I19000650001.

REFERENCES

- [1] M. Arjovsky, S. Chintala, and L. Bottou. Wasserstein generative adversarial networks. In *International Conference on Machine Learning*, pages 214–223, 2017. 2
- [2] C. Bundesen and A. Larsen. Visual transformation of size. *Journal of Experimental Psychology: Human Perception and Performance*, 1(3):214, 1975. 1
- [3] A. X. Chang, T. Funkhouser, L. Guibas, P. Hanrahan, Q. Huang, Z. Li, S. Savarese, M. Savva, S. Song, H. Su, et al. Shapenet: An information-rich 3d model repository. *arXiv preprint arXiv:1512.03012*, 2015. 5, 9
- [4] G. Chaurasia, S. Duchene, O. Sorkine-Hornung, and G. Drettakis. Depth synthesis and local warps for plausible image-based navigation. *ACM Transactions on Graphics (TOG)*, 32(3):30, 2013. 1
- [5] Q. Chen and V. Koltun. Photographic image synthesis with cascaded refinement networks. In *IEEE International Conference on Computer Vision (ICCV)*, volume 1, page 3, 2017. 5, 8
- [6] T. Chen, M.-M. Cheng, P. Tan, A. Shamir, and S.-M. Hu. Sketch2photo: Internet image montage. In *ACM transactions on graphics (TOG)*, volume 28, page 124. ACM, 2009. 1
- [7] X. Chen, J. Song, and O. Hilliges. Monocular neural image based rendering with continuous view control. In *Proceedings of the IEEE International Conference on Computer Vision*, 2019. 3
- [8] Y. Choi, M. Choi, M. Kim, J.-W. Ha, S. Kim, and J. Choo. Stargan: Unified generative adversarial networks for multi-domain image-to-image translation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8789–8797, 2018. 2
- [9] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. Imagenet: A large-scale hierarchical image database. 2009. 8
- [10] P. Esser, E. Sutter, and B. Ommer. A variational u-net for conditional appearance and shape generation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8857–8866, 2018. 1, 2, 6, 7, 8, 9
- [11] M. Everingham, L. Van Gool, C. K. Williams, J. Winn, and A. Zisserman. The pascal visual object classes (voc) challenge. *International journal of computer vision*, 88(2):303–338, 2010. 5
- [12] D. Feng, L. Rosenbaum, and K. Dietmayer. Towards safe autonomous driving: Capture uncertainty in the deep neural network for lidar 3d vehicle detection. In *2018 21st International Conference on Intelligent Transportation Systems (ITSC)*, pages 3266–3273. IEEE, 2018. 2
- [13] H. Gardner. *Frames of mind: The theory of multiple intelligences*. Hachette UK, 2011. 1
- [14] I. Goodfellow, Y. Bengio, and A. Courville. *Deep learning*. MIT press, 2016. 7
- [15] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. Generative adversarial nets. In *Advances in neural information processing systems*, pages 2672–2680, 2014. 1, 2
- [16] J. Hays and A. A. Efros. Scene completion using millions of photographs. *ACM Transactions on Graphics (TOG)*, 26(3):4, 2007. 1
- [17] K. He, G. Gkioxari, P. Dollár, and R. Girshick. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 2961–2969, 2017. 3
- [18] P. Hedman, T. Ritschel, G. Drettakis, and G. Brostow. Scalable inside-out image-based rendering. *ACM Transactions on Graphics (TOG)*, 35(6):231, 2016. 1
- [19] M. Heusel, H. Ramsauer, T. Unterthiner, B. Nessler, and S. Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In *Advances in Neural Information Processing Systems*, pages 6626–6637, 2017. 6, 8, 9
- [20] X. Huang, X. Cheng, Q. Geng, B. Cao, D. Zhou, P. Wang, Y. Lin, and R. Yang. The apollo-scape dataset for autonomous driving. In *Proceedings of the IEEE Conference on CVPR Workshops*, pages 954–960, 2018. 3
- [21] P. Isola and C. Liu. Scene collaging: Analysis and synthesis of natural images with semantic layers. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3048–3055, 2013. 1
- [22] P. Isola, J.-Y. Zhu, T. Zhou, and A. A. Efros. Image-to-image translation with conditional adversarial networks. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5967–5976. IEEE, 2017. 2, 5
- [23] J. Johnson, A. Alahi, and L. Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *European conference on computer vision*, pages 694–711. Springer, 2016. 5
- [24] T. Karras, T. Aila, S. Laine, and J. Lehtinen. Progressive growing of gans for improved quality, stability, and variation. *International Conference on Learning Representations (ICLR)*, 2018. 2
- [25] N. Kholgade, T. Simon, A. Efros, and Y. Sheikh. 3d object manipulation in a single photograph using stock 3d models. *ACM Transactions on Graphics (TOG)*, 33(4):127, 2014. 3
- [26] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations (ICLR)*, 2015. 6
- [27] D. P. Kingma and M. Welling. Auto-encoding variational bayes. In *International Conference on Learning Representations (ICLR)*, 2014. 1, 2, 6
- [28] J.-F. Lalonde, D. Hoiem, A. A. Efros, C. Rother, J. Winn, and A. Criminisi. Photo clip art. In *ACM transactions on graphics (TOG)*, volume 26, page 3. ACM, 2007. 1
- [29] J. J. Lim, H. Pirsiavash, and A. Torralba. Parsing IKEA Objects: Fine Pose Estimation. *ICCV*, 2013. 3
- [30] X. Liu, W. Liu, H. Ma, and H. Fu. Large-scale vehicle re-identification in urban surveillance videos. In *2016 IEEE International Conference on Multimedia and Expo (ICME)*, pages 1–6. IEEE, 2016. 2
- [31] Z. Liu, P. Luo, S. Qiu, X. Wang, and X. Tang. Deepfashion: Powering robust clothes recognition and retrieval with rich annotations. In *Proceedings of the IEEE conference on CVPR*, pages 1096–1104, 2016. 2
- [32] M. Lucic, K. Kurach, M. Michalski, S. Gelly, and O. Bousquet. Are gans created equal? a large-scale study. In *Advances in neural information processing systems*, pages 698–707, 2018. 8
- [33] L. Ma, X. Jia, Q. Sun, B. Schiele, T. Tuytelaars, and L. Van Gool. Pose guided person image generation. In *Advances in Neural Information Processing Systems*, pages 406–416, 2017. 2
- [34] P. A. Marín-Reyes, L. Bergamini, J. Lorenzo-Navarro, A. Palazzi, S. Calderara, and R. Cucchiara. Unsupervised vehicle re-identification using triplet networks. In *2018 IEEE/CVF Conference on CVPR Workshops (CVPRW)*, pages 166–1665. IEEE, 2018. 2
- [35] M. Mirza and S. Osindero. Conditional generative adversarial nets. *arXiv preprint arXiv:1411.1784*, 2014. 2
- [36] T. Nguyen-Phuoc, C. Li, L. Theis, C. Richardt, and Y.-L. Yang. Hologan: Unsupervised learning of 3d representations from natural images. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 7588–7597, 2019. 3
- [37] M. Oechsle, L. Mescheder, M. Niemeyer, T. Strauss, and A. Geiger. Texture fields: Learning texture representations in function space. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 4531–4540, 2019. 3
- [38] K. Olszewski, S. Tulyakov, O. Woodford, H. Li, and L. Luo. Transformable bottleneck networks. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 7648–7657, 2019. 3
- [39] R. Ortiz-Cayon, A. Djelouah, and G. Drettakis. A bayesian approach for selective image-based rendering using superpixels. In *International Conference on 3D Vision-3DV*, 2015. 1
- [40] R. Ortiz-Cayon, A. Djelouah, F. Massa, M. Aubry, and G. Drettakis. Automatic 3d car model alignment for mixed image-based rendering. In *2016 Fourth International Conference on 3D Vision (3DV)*, pages 286–295. IEEE, 2016. 1
- [41] A. Palazzi, L. Bergamini, S. Calderara, and R. Cucchiara. End-to-end 6-dof object pose estimation through differentiable rasterization. In *Second Workshop on 3D Reconstruction Meets Semantics (3DRMS)*, 2018. 3
- [42] E. Park, J. Yang, E. Yumer, D. Ceylan, and A. C. Berg. Transformation-grounded image generation network for novel 3d view synthesis. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 702–711. IEEE, 2017. 1, 2, 6, 7, 8, 9
- [43] A. Paszke, S. Gross, S. Chintala, G. Chanan, E. Yang, Z. DeVito, Z. Lin, A. Desmaison, L. Antiga, and A. Lerer. Automatic differentiation in pytorch. 2017. 6
- [44] G. Pavlakos, X. Zhou, A. Chan, K. G. Derpanis, and K. Daniilidis. 6-dof object pose from semantic keypoints. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*, pages 2011–2018. IEEE, 2017. 3, 6
- [45] X. Qi, Q. Chen, J. Jia, and V. Koltun. Semi-parametric image synthesis. In *Proceedings of the IEEE Conference on CVPR*, pages 8808–8816, 2018. 1, 3, 5, 8
- [46] A. Radford, L. Metz, and S. Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. *International Conference on Learning Representations (ICLR)*, 2016. 2
- [47] K. Rematas, C. H. Nguyen, T. Ritschel, M. Fritz, and T. Tuytelaars. Novel views of objects from a single image. *IEEE transactions on pattern analysis and machine intelligence*, 39(8):1576–1590, 2017. 3
- [48] O. Ronneberger, P. Fischer, and T. Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015. 6
- [49] Y. Shen, T. Xiao, H. Li, S. Yi, and X. Wang. Learning deep neural networks for vehicle re-id with visual-spatio-temporal path proposals. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1900–1909, 2017. 2
- [50] R. N. Shepard and J. Metzler. Mental rotation of three-dimensional objects. *Science*, 171(3972):701–703, 1971. 1
- [51] P. Shilane, P. Min, M. Kazhdan, and T. Funkhouser. The princeton shape

- benchmark. In *Proceedings Shape Modeling Applications, 2004.*, pages 167–178. IEEE, 2004. 5
- [52] A. Siarohin, E. Sangineto, S. Lathuilière, and N. Sebe. Deformable gans for pose-based human image generation. In *CVPR 2018-Computer Vision and Pattern Recognition*, 2018. 2
- [53] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. 5
- [54] A. Sinha, A. Unmesh, Q. Huang, and K. Ramani. Surfnet: Generating 3d shape surfaces using deep residual networks. In *IEEE CVPR*, 2017. 3
- [55] V. Sitzmann, J. Thies, F. Heide, M. Nießner, G. Wetzstein, and M. Zollhofer. Deepvoxels: Learning persistent 3d feature embeddings. In *Proceedings of the IEEE Conference on CVPR*, pages 2437–2446, 2019. 3
- [56] S.-H. Sun, M. Huh, Y.-H. Liao, N. Zhang, and J. J. Lim. Multi-view to novel view: Synthesizing novel views with self-learned confidence. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 155–171, 2018. 2, 6, 7, 8, 9
- [57] X. Sun, J. Wu, X. Zhang, Z. Zhang, C. Zhang, T. Xue, J. B. Tenenbaum, and W. T. Freeman. Pix3d: Dataset and methods for single-image 3d shape modeling. In *Proceedings of the IEEE Conference on CVPR*, pages 2974–2983, 2018. 5
- [58] M. Tatarchenko, A. Dosovitskiy, and T. Brox. Multi-view 3d models from single images with a convolutional network. In *European Conference on Computer Vision*, pages 322–337. Springer, 2016. 1, 2, 6, 7, 8
- [59] M. Teichmann, M. Weber, M. Zoellner, R. Cipolla, and R. Urtasun. Multinet: Real-time joint semantic reasoning for autonomous driving. In *2018 IEEE Intelligent Vehicles Symposium (IV)*, pages 1013–1020. IEEE, 2018. 2
- [60] S. Tulsiani and J. Malik. Viewpoints and keypoints. In *Proceedings of the IEEE Conference on CVPR*, pages 1510–1519, 2015. 3, 6
- [61] T.-C. Wang, M.-Y. Liu, J.-Y. Zhu, A. Tao, J. Kautz, and B. Catanzaro. High-resolution image synthesis and semantic manipulation with conditional gans. In *Proceedings of the IEEE Conference on CVPR*, pages 8798–8807, 2018. 2
- [62] Z. Wang, L. Tang, X. Liu, Z. Yao, S. Yi, J. Shao, J. Yan, S. Wang, H. Li, and X. Wang. Orientation invariant feature embedding and spatial temporal regularization for vehicle re-identification. In *The IEEE International Conference on Computer Vision (ICCV)*, Oct 2017. 3
- [63] J. Wu, Y. Wang, T. Xue, X. Sun, B. Freeman, and J. Tenenbaum. Marrnet: 3d shape reconstruction via 2.5 d sketches. In *Advances in neural information processing systems*, pages 540–550, 2017. 3, 5
- [64] J. Wu, T. Xue, J. J. Lim, Y. Tian, J. B. Tenenbaum, A. Torralba, and W. T. Freeman. 3d interpreter networks for viewer-centered wireframe modeling. *International Journal of Computer Vision*, 126(9):1009–1026, 2018. 3
- [65] W. Xian, P. Sangkloy, V. Agrawal, A. Raj, J. Lu, C. Fang, F. Yu, and J. Hays. Texturegan: Controlling deep image synthesis with texture patches. In *Proceedings of the IEEE Conference on CVPR*, pages 8456–8465, 2018. 6
- [66] Y. Xiang, W. Kim, W. Chen, J. Ji, C. Choy, H. Su, R. Mottaghi, L. Guibas, and S. Savarese. Objectnet3d: A large scale database for 3d object recognition. In *European Conference on Computer Vision*, pages 160–176. Springer, 2016. 3
- [67] Y. Xiang, R. Mottaghi, and S. Savarese. Beyond pascal: A benchmark for 3d object detection in the wild. In *Applications of Computer Vision (WACV), 2014 IEEE Winter Conference on*, pages 75–82. IEEE, 2014. 3, 4, 5, 6, 10
- [68] F. Xiao, H. Liu, and Y. J. Lee. Identity from here, pose from there: Self-supervised disentanglement and generation of objects using unlabeled videos. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 7013–7022, 2019. 3
- [69] J. Yang, S. E. Reed, M.-H. Yang, and H. Lee. Weakly-supervised disentangling with recurrent transformations for 3d view synthesis. In *Advances in Neural Information Processing Systems*, pages 1099–1107, 2015. 1, 2
- [70] M. Yang, K. Yu, C. Zhang, Z. Li, and K. Yang. Denseaspp for semantic segmentation in street scenes. In *Proceedings of the IEEE Conference on CVPR*, pages 3684–3692, 2018. 2
- [71] J. Yu, Z. Lin, J. Yang, X. Shen, X. Lu, and T. S. Huang. Generative image inpainting with contextual attention. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018. 2
- [72] H. Zhang, T. Xu, H. Li, S. Zhang, X. Wang, X. Huang, and D. N. Metaxas. Stackgan: Text to photo-realistic image synthesis with stacked generative adversarial networks. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 5907–5915, 2017. 2
- [73] B. Zhao, X. Wu, Z.-Q. Cheng, H. Liu, Z. Jie, and J. Feng. Multi-view image generation from a single-view. In *2018 ACM Multimedia Conference on Multimedia Conference*, pages 383–391. ACM, 2018. 2
- [74] L. Zheng, L. Shen, L. Tian, S. Wang, J. Wang, and Q. Tian. Scalable person re-identification: A benchmark. In *Computer Vision, IEEE*

- International Conference on*, 2015. 2
- [75] Q.-Y. Zhou, J. Park, and V. Koltun. Open3D: A modern library for 3D data processing. *arXiv:1801.09847*, 2018. 6
- [76] T. Zhou, S. Tulsiani, W. Sun, J. Malik, and A. A. Efros. View synthesis by appearance flow. In *European conference on computer vision*, pages 286–301. Springer, 2016. 1, 2
- [77] J.-Y. Zhu, T. Park, P. Isola, and A. A. Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. *IEEE International Conference on Computer Vision (ICCV)*, 2017. 2, 5
- [78] J.-Y. Zhu, R. Zhang, D. Pathak, T. Darrell, A. A. Efros, O. Wang, and E. Shechtman. Toward multimodal image-to-image translation. In *Advances in Neural Information Processing Systems*, pages 465–476, 2017. 2
- [79] J.-Y. Zhu, Z. Zhang, C. Zhang, J. Wu, A. Torralba, J. Tenenbaum, and B. Freeman. Visual object networks: image generation with disentangled 3d representations. In *Advances in Neural Information Processing Systems*, pages 118–129, 2018. 1, 2, 3, 5, 6, 7, 8, 9



Andrea Palazzi MSc in computer engineering in 2015, PhD in 2019 at University of Modena and Reggio Emilia, Italy. His PhD research revolved around the topics of driver's attention prediction and gaze estimation, use of synthetic data to train deep neural networks, synthesis of novel views from single monocular images. He is currently Senior Deep Learning Engineer at NomiTri, where he tackles the challenge of bringing visual intelligence to smartphones and edge devices.



Luca Bergamini is currently in his last year of PhD in the AlmageLab at the University of Modena and Reggio Emilia, where he received his master degree in 2017. His research interests lie in the area of computer vision and deep learning. He has recently joined Lyft as a research SWE.



Simone Calderara received a computer engineering master's degree in 2005 and the PhD degree in 2009 from the University of Modena and Reggio Emilia, where he is currently an assistant professor within the Imagelab group. His current research interests include computer vision and machine learning applied to human behavior analysis, visual tracking in crowded scenarios, and time series analysis for forensic applications. He is a member of the IEEE.



Rita Cucchiara (PhD in Computer Engineering 1992, University of Bologna, Italy) is full professor at the University of Modena and Reggio Emilia, Italy, where she heads the AlmageLab group. Director of the Italian Lab CINI Of Artificial Intelligence and intelligent Systems. Former President of the Italian CVPL association. IAPR Fellow since 2006, ELLIS fellow since 2019. Part of the Governing Board of IAPR and a member of the the Advisory Board of CVF. Area, Program, General Chair of several top-tier conferences. She published more than 400 papers on pattern recognition, computer vision and multimedia.