

# Adversarial Inverse Graphics Networks: Learning 2D-to-3D Lifting and Image-to-Image Translation from Unpaired Supervision

Hsiao-Yu Fish Tung\* Adam W. Harley\* William Seto\* Katerina Fragkiadaki  
Carnegie Mellon University  
{htung, aharley, wseto, katef}@cs.cmu.edu

## Abstract

Researchers have developed excellent feed-forward models that learn to map images to desired outputs, such as to the images' latent factors, or to other images, using supervised learning. Learning such mappings from unlabelled data, or improving upon supervised models by exploiting unlabelled data, remains elusive. We argue that there are two important parts to learning without annotations: (i) matching the predictions to the input observations, and (ii) matching the predictions to known priors. We propose Adversarial Inverse Graphics networks (AIGNs): weakly supervised neural network models that combine feedback from rendering their predictions, with distribution matching between their predictions and a collection of ground-truth factors. We apply AIGNs to 3D human pose estimation and 3D structure and egomotion estimation, and outperform models supervised by only paired annotations. We further apply AIGNs to facial image transformation using super-resolution and inpainting renderers, while deliberately adding biases in the ground-truth datasets. Our model seamlessly incorporates such biases, rendering input faces towards young, old, feminine, masculine or Tom Cruise-like equivalents (depending on the chosen bias), or adding lip and nose augmentations while inpainting concealed lips and noses.

## 1. Introduction

Humans imagine far more than they see. In Figure 1 (b), we imagine the hidden arms and legs of the sitting woman. In Figure 1 (c), we imagine forward motion of the camera, as opposed to the road drifting backwards underneath the camera. We arrive at these interpretations – *i.e.*, predictions of latent factors – by referring to *priors* on how the world works.

This work proposes Adversarial Inverse Graphics Networks (AIGNs), a model that learns to map images to la-

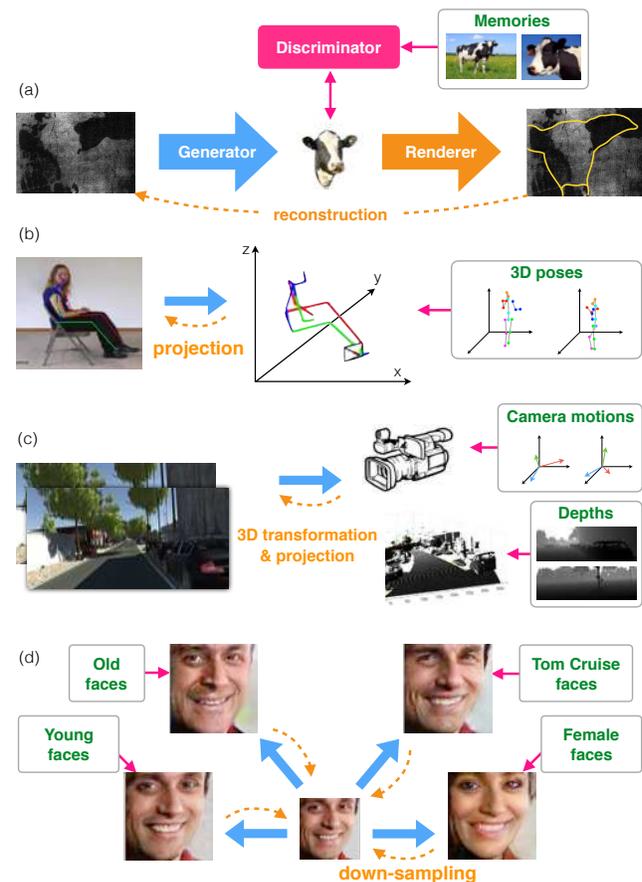


Figure 1. **Adversarial Inverse Graphics Networks** combine feedback from differentiable rendering of the predictions, with priors, imposed through distribution matching between predictions and “memories”. The general architecture is shown (a) with Dallenbach’s classic cow illusion [4], and subsequently with our experimental tasks: (b) 3D human pose estimation, (c) extraction of 3D depth and camera motion from a pair of frames, and (d) creative and controllable image generation. Each task uses a unique domain-specific renderer, depicted in orange dashed arrows.

tent factors using feedback from the rendering of its predictions, as well as distribution matching between the predic-

\*equal contribution

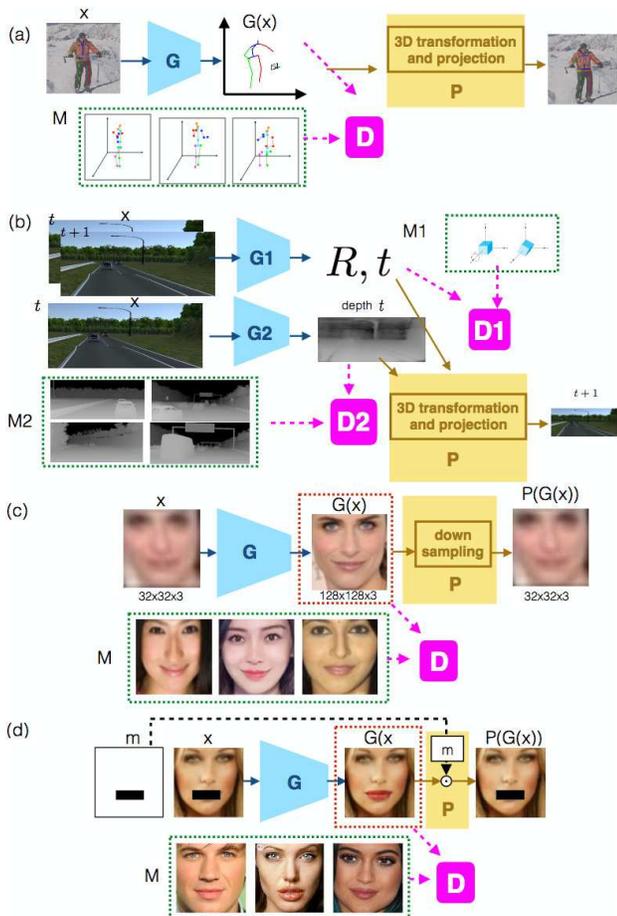


Figure 2. **Adversarial inverse graphics architectures** for (a) 3D human pose estimation, (b) Structure from motion, (c) super-resolution, and (d) inpainting.

tions and a stored collection of ground truth latent factors. The ground-truth collection does not need to be directly related to the current input – it can be a disordered set of labels. We call this *unpaired supervision*. The renderers employed are differentiable, parameter-free, and task specific, e.g., camera projection, camera motion, downsampling, and masking. Figure 2 depicts architectures for AIGNs for the tasks of structure from motion, 3D human pose estimation, super-resolution, and in-painting.

When we (deliberately) *bias the ground truth* collection of an AIGN to reflect a distorted reality, surprising renders arise. For super-resolution, AIGNs can make people look older, younger, more feminine, more masculine or more like Tom Cruise, by simply curating the discriminators’ ground truth to include images of old, young, female, male, or Tom Cruise pictures, respectively. For inpainting, AIGNs can make people appear with bigger lips or bigger noses, again by curating the discriminators’ ground truth to include faces with big lips or big noses. These observations inspire a compelling analogy: the AIGN’s ground truth collection is

its *memory*, and the renders are its *imagination*. When the memories are biased, the imaginations reflect a distorted reality.

Our model is related to, and builds upon, many recent works in the literature. Inverse-graphics networks [17] use *parametric* deconvolutional networks and strong supervision (*i.e.*, annotation of images with their ground truth imaginations) for decomposing an image into interpretable factors, e.g., albedo and shading. Our model instead employs *parameter-free* renderers, and makes use of unpaired weak supervision. Similar to our approach, 3D interpreter networks [28] use a reconstruction loss in the form of a 2D re-projection error of predicted 3D keypoints, along with paired supervised training for 3D human pose estimation. Our model complements the reconstruction loss with adversarial losses on the predicted 3D human poses, and shows performance superior to the 3D interpreter. Conditional Generative Adversarial Networks (GANs) have used a combination of adversarial and L2 losses, e.g., for inpainting [20], 3D voxel generation [30], super-resolution [5], and image-to-image translation [14]. In these models, the adversarial loss is used to avert the regression-to-mean problem in standard L2 regression (*i.e.* blurring). Such models are feed-forward; they do not have a renderer and reconstruction feedback. As a result, they can only exploit supervision in the form of annotated pairs e.g., of an RGB image and its corresponding 3D voxel grid [30]. Our model extends such supervised conditional GAN formulations through self-consistency via inverse rendering, and matching between *rendered predictions and the visual input*, rather than between predictions and the ground truth. This feedback loop allows weakly supervised distribution matching to work (using *unpaired* annotations for the discriminators), removing the need for direct label matching.

As unsupervised models, AIGNs do not discriminate between training and test phases. Extensive experiments in a variety of tasks show that their performance consistently improves over their supervised alternatives, by adapting in a self-supervised manner to the statistics of the test set.

## 2. Adversarial Inverse Graphics Networks

AIGN architectures for various tasks are shown in Figure 2. Given an image or image pair  $x$ , generator networks  $G_i$ ,  $i \in 1..K$  map  $x$  to a set of predictions  $\mathcal{G}(x) = \{G_i(x), i \in 1..K\}$ . Then, a task-specific differentiable renderer  $P(\mathcal{G}(x))$  renders predictions back to the original input space. Discriminator networks  $D_i$ ,  $i \in 1..K$  are trained to discriminate between predictions  $G_i(x)$  and true *memories* of appropriate form  $M_i$ . Discriminators have to assign large probability to true memories and low probability to generators’ predictions. Given a set of images (or image pairs depending on the task)  $X = \{x_1, x_2, \dots, x_n\}$ , generators are trained to minimize L2 distance between rendered imagina-

tions and input  $x$  and simultaneously maximize discriminators' confusion. Our loss then reads:

$$\min_{\mathcal{G}} \max_{\mathcal{D}} \mathbb{E}_{x \in X} \underbrace{\|P(\mathcal{G}(x)) - x\|_2}_{\text{reconstruction loss}} + \underbrace{\beta \sum_{i=1}^K \log D_i(M_i) + \log(1 - D_i(G_i(x)))}_{\text{adversarial loss}}, \quad (1)$$

where  $\mathcal{D}$  denotes the set of discriminator networks and  $\mathcal{G}$  the set of generator networks and  $\beta$  the relative weight of reconstruction and adversarial losses. Since paired ground-truth is not used anywhere, both reconstruction and adversarial losses can be used both at train and test time. Our model can also benefit from strong supervision through paired annotations -pairs of visual input with desired predictions- at training time for training the generator networks. We will use the term *adversarial priors* to denote adversarial losses over latent predicted factors of our model.

In the next sections, we present applications of AIGNs to the tasks of (i) 3D human pose estimation, (ii) extraction of 3D depth and egomotion from a pair of frames, which we call Structure from Motion (SfM), (iii) super-resolution, and (iv) inpainting. For application of AIGNs to intrinsic image decomposition, please see our earlier unpublished manuscript [26].

### 2.1. 3D Human Pose Estimation

The AIGN architecture for 3D human pose estimation is depicted in Figure 2 (a). Given an image crop centered around a person detection, the task is to predict a 3D skeleton for the depicted person. We decompose the 3D human shape into a view-aligned 3D linear basis model and a rotation matrix:  $x_{3D} = R \cdot \sum_{j=1}^{|\mathcal{B}|} \alpha_j B_j$ . Our shape basis  $\mathcal{B}$  is obtained using PCA on orientation-aligned 3D poses in our training set, where orientation is measured by the direction of the normal to the hip vector. We keep 60 components out of a total of 96 (three coordinates of 32 keypoints). The dimensionality reduction is small, and indeed, we just use bases weights for ease of prediction, relying on our adversarial priors (rather than PCA) to regularize the 3D shape prediction.

Our generator network first predicts a set of 2D bodyjoint heatmaps and then maps those to basis weights  $\mathbf{w} \in \mathbb{R}^{60}$ , focal length  $f$ , and Euler rotation angles  $\alpha, \beta, \gamma$  so that the 3D rotation of the shape reads  $R = R^x(\alpha)R^y(\beta)R_t^z(\gamma)$ , where  $R^x(\theta)$  denotes rotation around the x-axis by angle  $\theta$ . Our renderer then projects the 3D keypoints  $x_{3D}$  to 2D keypoints  $x_{2D}^{proj}$ , all in homogeneous coordinates:

$$x_{2D}^{proj} = P \cdot x_{3D} + \begin{bmatrix} c_x \\ c_y \\ 0 \end{bmatrix}, \quad (2)$$

where  $P = \begin{bmatrix} f & 0 & 0 & 0 \\ 0 & f & 0 & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix}$  is the camera projection ma-

trix. The reconstruction loss is the  $L2$  norm between the reprojected coordinates  $x_{2D}^{proj}$  and 2D coordinates obtained by the argmax of predicted 2D heatmaps. A discriminator network discriminates between our generated keypoints  $x_{3D}$  and a database of 3D human skeletons, which does not contain the ground-truth (paired) 3D skeleton for the input image, but rather previous poses (*i.e.*, “memories”).

### 2.2. Structure from Motion

Simultaneous Localization And Mapping (SLAM) methods have shown impressive results on estimating camera pose and 3D point clouds from monocular, stereo, or RGB-D video sequences [22, 16] using alternate optimization over camera poses and the 3D scene pointcloud. There has been a recent interest in integrating learning to aid geometric approaches handle moving objects and low-texture image regions. The work of Handa et al. [11] contributed a deep learning library with many geometric operations including a differentiable camera projection layer, similar to those used in SfM networks [21, 33], 3D image interpreters [29], and depth from stereo CNN [9]. Our SfM AIGN architecture, depicted in Figure 2 (b), build upon those works. Given a pair of consecutive frames, the task is to predict the relative camera motion and a depth map for the first frame.

We have two generator networks: an egomotion network and a depth network. The egomotion network takes the following inputs: a concatenation of two consecutive frames  $I_1, I_2$ , their 2D optical flow estimated using our implementation of FlowNet [6], the angles of the optical flow, and a static angle field. The angle field denotes each pixel's angle from the camera's principal point (in radians), while the flow angles denote the angular component of the optical flow. The egomotion network produces the camera's 3D relative rotation and translation  $\{R, T\} \in SE(3)$  as output. We represent 3D camera relative rotation using an Euler angle representation:  $R = R^x(\alpha)R^y(\beta)R_t^z(\gamma)$  where  $R^x(\theta)$  denotes rotation around the x-axis by the angle  $\theta$ . The egomotion network is a standard convolutional architecture (similar to VGG16 [23]), but following the last convolution layer, there is a single fully-connected layer producing a 6D vector, representing the Euler angles of the rotation matrix, and the translation vector.

The depth network takes as input the first frame  $I_1$  and estimates 3D scene log depth at every pixel,  $\{d\} \in \mathbb{R}^{w \times h}$ . The architecture of this network is the same as that of the egomotion network, except deconvolution layers (with skip connections) replace the fully-connected layer, producing a depth estimate at every pixel. Given generated depth  $d$  for the first frame  $I_1$  and known camera intrinsics – focal length  $f$  and principal point  $(c_x, c_y)$  – we obtain the correspond-

ing 3D point cloud for  $I_1$ :  $\mathbf{X}_1^i = \begin{bmatrix} X_1^i \\ Y_1^i \\ Z_1^i \end{bmatrix} = \frac{d^i}{f} \begin{bmatrix} x_1^i - c_x \\ y_1^i - c_y \\ f \end{bmatrix}$ ,

where  $(x_1^i, y_1^i)$  are the column and row coordinates and  $d^i$  the predicted depth of the  $i^{\text{th}}$  pixel in frame  $I_1$ . Our renderer in this task transforms the point cloud  $\mathbf{X}_1$  according to the estimated camera motion:  $\mathbf{X}_2 = R \cdot \mathbf{X}_1 + T$  (similar to SE3-Nets [3] and the recent SfM-Nets [21, 33]), and projects the 3D points back to image pixels  $x_2, y_2$  using camera projection.

The optical flow vector of pixel  $i$  is then  $(U^i, V^i) = (x_2^i - x_1^i, y_2^i - y_1^i)$ . Our reconstruction loss corresponds to the well known brightness constancy principle, that pixel appearance does not change under small pixel motion:

$$\mathcal{L}^{\text{photo}} = \frac{1}{w h} \sum_{x,y} \|I_1(x, y) - I_2(x+U(x, y), y+V(x, y))\|_1 \quad (3)$$

We have two discriminator networks:  $D_1$  manages the realism of predicted camera motions by enforcing physics constraints in a statistical manner: when the camera is fixed to the top of a car, the roll and pitch angles are close to zero due to the nature of the car motion. This network has three fully-connected layers (of size 128, 128, and 64), and discriminates between real and fake relative transformation matrices. The second discriminator,  $D_2$ , manages the realism of 3D depth predictions, using a collection of ground-truth depth maps from unrelated sequences. This discriminator is *fully convolutional*. That is, it has one probabilistic output at each image sub-region, as we are interested in the realism of local depth texture. We found the depth discriminator to effectively handle the scale ambiguity of monocular 3D reconstruction: in a monocular setup without priors, the re-projection error cannot tell if an object is far and big, or near and small. The depth discriminator effectively imposes a prior regarding world scale, as well as depth realism on fine-grained details.

### 2.3. Super-resolution

Deep neural network architectures have recently shown excellent performance on the task of super-resolution [5, 2] and non-blind inpainting [20, 32] (inpainting where the mask is always at the same part of the image). Adversarial losses have been used to resolve the regression to the mean problem of standard  $L_2$  loss [18, 20]. Our model unifies recent works of [24] and [32] that combine adversarial and reconstruction losses for super-resolution and inpainting respectively, without paired supervision, same as our model. However, weak (unpaired) supervision might be unnecessary for super-resolution or inpainting as an unlimited amount of ground-truth pairs can be easily collected by downsampling or masking images. Our AIGN focuses instead on *biased* super-resolution and inpainting, as a tool for creative and controllable image generation.

The AIGN architecture for super-resolution is depicted in Figure 2 (c). Our generator network takes as input a low-resolution image  $x \in \mathbb{R}^{w \times h}$  and produces a high resolution image  $G(x) \in \mathbb{R}^{4W \times 4H}$  after a number of residual neural blocks [12]. Our renderer  $P(G(x))$  is a downsampler that reduces the size of the output image by four times. We implement it using average pooling with the appropriate stride. Our reconstruction loss is the  $L_2$  distance *between the input image and the rendered imagination*  $P(G(x))$ . Our discriminator takes as input high resolution predicted images, as well as memories of high resolution images, unrelated to the current input. Thus far, our model is similar to the concurrent work Sønderby *et al.* [24].

Unsupervised super-resolution networks [24] may not be necessary, since large supervised training datasets of low- and high-resolution image pairs can be collected by Gaussian blurring and downsampling. We instead focus on what we call *biased super-resolution* for face images. We bias our discriminator ground-truth memories to contain high resolution images of a particular image category, *e.g.*, females, males, young, old faces, or faces of a particular individual. AIGNs then mix the low-frequencies of the input image (preserved via our reconstruction loss) with high frequencies from the memories (imposed by the adversarial losses), the relative weight  $\beta$  between reconstruction and adversarial loss in Eq. 1 controls such mixing. The result is realistically looking faces whose age, gender or identity has been transformed, as shown in Figure 5. Such facial transformations are completely unsupervised; the model has never seen a pair of the same person old and young (or male and female).

### 2.4. Inpainting

The AIGN architecture for inpainting is depicted in Figure 2 (d). The input is a “masked” image  $x$ , that is, an image whose content is covered by a black contiguous mask  $m$ . Our generator produces a complete (inpainted) image  $G(x)$ . The rendering function  $P$  in this case is a masking operation:  $P(G(x), m) = m \odot G(x)$ , where  $\odot$  denotes pointwise multiplication. Our discriminator receives inpainted imaginations  $G(x)$  and memories of complete face images  $M$ , unrelated to our current input images. Our model is trained then to predict complete, inpainted images that when masked will match the input image  $x$ . Thus far, our model is similar to [32].

Unsupervised inpainting networks [32] may not be necessary, since large supervised training datasets of paired masked and complete images can be collected via image masking. We instead focus on *biased inpainting* of face images. We bias the discriminator’s ground-truth memories to contain complete images with a particular desired characteristic, in the location of the mask  $m$ . For example, if the mask covers the mouth of a person, we can bias dis-

	Direct	Discuss Eat	Greet	Phone Photo	Pose	Purchase	Sit	SitDown	Smoke	Wait	Walk	Average		
Forward2Dto3D	75.2	118.4	165.7	95.9	149.1	154.1	77.7	176.9	186.5	193.7	142.7	99.8	74.7	128.9
3Dinterpr [28]	56.3	77.5	96.2	71.6	96.3	106.7	59.1	109.2	111.9	<b>111.9</b>	124.2	93.3	58.0	88.6
Monocap [34]	78.0	78.9	88.1	93.9	102.1	115.7	71.0	<b>90.6</b>	121.0	118.2	102.5	82.6	75.62	92.3
AIGN (ours)	<b>53.7</b>	<b>71.5</b>	<b>82.3</b>	<b>58.6</b>	<b>86.9</b>	<b>98.4</b>	<b>57.6</b>	<b>104.2</b>	<b>100.0</b>	112.5	<b>83.3</b>	<b>68.9</b>	<b>57.0</b>	<b>79.0</b>

Table 1. **3D reconstruction error** in H3.6M using ground-truth 2D keypoints as input.

	Direct	Discuss Eat	Greet	Phone Photo	Pose	Purchase	Sit	SitDown	Smoke	Wait	Walk	Average		
Forward2Dto3D	80.2	92.4	102.8	92.5	115.5	79.9	119.5	136.7	136.7	144.4	109.3	94.2	80.2	104.6
3Dinterpr [28]	78.6	<b>90.8</b>	92.5	89.4	108.9	112.4	77.1	<b>106.7</b>	127.4	139.0	103.4	91.4	79.1	98.4
AIGN (ours)	<b>77.6</b>	91.4	<b>89.9</b>	<b>88</b>	<b>107.3</b>	<b>110.1</b>	<b>75.9</b>	107.5	<b>124.2</b>	<b>137.8</b>	<b>102.2</b>	<b>90.3</b>	<b>78.6</b>	<b>97.2</b>

Table 2. **3D reconstruction error** in H3.6M using detected 2D keypoints as input.

criminator’s ground-truth memories to only contain people with big lips. In this case, our generator will produce inpainted images  $G(x)$ , that have this localized characteristic in a smooth photorealistic blend, in order to confuse the discriminator  $D$ .

For further details on the proposed architectures please see the supplementary material.

### 3. Experiments

#### 3.1. 3D human pose estimation from a static image

We use the Human3.6M (H3.6M) dataset of Ionescu *et al.* [13], the largest available dataset with annotated 3D human poses. This dataset contains videos of actors performing activities and provides annotations of body joint locations in 2D and 3D at every frame, recorded from a Vicon system. We split the videos by the human subjects, with five subjects (S1, S5, S6, S7, S8) for training and two subjects (S9, S11) for testing, following the split of previous works [34]. For both sets, we use one third of the original frame rate.

We consider a variety of supervised setups and baselines for our 3D human pose predictor, which we detail below. We first train our network using synthetic data augmentation (Sup 1), following the protocol of Wu *et al.* [28]: A 3D example skeleton is first perturbed, a 3D rotation  $R$  and focal length  $f$  are sampled, and the resulting rotated shape is projected to 2D. We further train our network using real paired 2D-to-3D training data from H3.6M (Sup 2). Our generator network trained with Sup1+Sup2 we will call it *Forward2Dto3D* net, as it resembles a standard supervised model for 3D human pose estimation. We further finetune using a reconstruction loss (2D reprojection error) in the test data (Sup 3). Our generator network trained with Sup1+Sup2+Sup3 we will call it *3D interpreter* due to its clear similarity with Wu *et al.* [28]. Since the original source code is not available, we re-implement it and use it as one of our baselines. Our AIGN model,

along with Sup1+Sup2+Sup3, uses unsupervised adversarial losses in the test data using randomly selected 3D training poses (Sup4). We compare AIGN to 3D interpreter and Forward2Dto3D baselines in two setups for 3D lifting: (a) using ground-truth 2D body joints provided by H3.6M as input, and (b) using 2D body joints provided by the state-of-the-art Convolutional Pose Machine detector [27]. We used an SVM regressor to map keypoint definitions of Wei *et al.* [27] to the one of H3.6M dataset. When using ground-truth 2D body joints as input we also compare against the publicly available 3D pose code for MonoCap [34], an optimization method that uses a sparsity prior across an over-complete dictionary of 3D poses, and minimizes the reprojection error via Expectation-Maximization. We consider one image as input for all the models for a fair comparison (MonoCap was originally proposed assuming a video sequence as input).

**Evaluation metrics.** Given a set of estimated 3D joint locations  $\hat{\mathbf{x}}_1 \cdots \hat{\mathbf{x}}_K$  and corresponding ground-truth 3D joint locations  $\mathbf{x}_1^* \cdots \mathbf{x}_K^*$ , the reconstruction error is defined as the 3D per-joint error after the torsos are aligned to face the front with transformation  $\mathcal{T} : \frac{1}{K} \sum_{i=1}^K \|\mathcal{T}(\hat{\mathbf{x}}_i) - \mathcal{T}(\mathbf{x}_i^*)\|$ . We show the 3D reconstruction error (in millimeters) of our model and baselines in Tables 1 and 2, organized according to activity, following the presentation format of Zhou *et al.* [34], though *only one model was used across all activities* for our method and baselines (for MonoCap this means using the same dictionary for optimization in all images).

The AIGN outperforms the baselines, especially for ground-truth keypoints. This suggests it would be valuable to finetune 2D keypoint detector features as well, instead of keeping them fixed. Adversarial priors allow the model not to diverge when finetuned on new (unlabelled) data, as they ensure anthropomorphism and plausibility of the detected poses. For additional 3D pose results, please see the supplementary material.

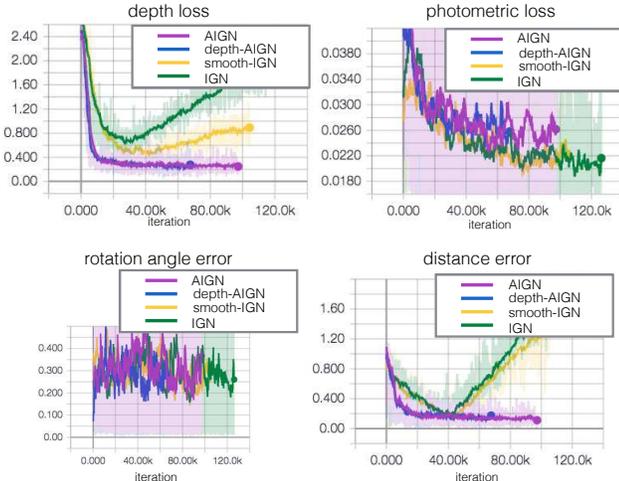


Figure 3. Depth, camera motion (relative rotation and translation) and photometric reprojection error curves during training of (a) the AIGN model with adversarial losses on both depth and camera motion (*AIGN*), (b) an AIGN with adversarial losses only on depth (*depth-AIGN*), (c) a model with reconstruction loss and depth smoothness loss (*smooth-IGN*) and (d) a model with a reconstruction loss only (*IGN*). Adversarial priors handle scale ambiguity of reconstruction loss, and thus models do not diverge.

### 3.2. Structure from Motion

We use Virtual KITTI (VKITTI) [8], a synthetic dataset that depicts videos taken from a camera mounted on a car driving in an urban center. This dataset contains scenes and camera viewpoints similar to the real KITTI dataset [10]. We use the VKITTI dataset rather than KITTI, because the real-world dataset has large errors in the ground truth camera motion for short sub-sequences, whereas the synthetic ground truth is precise. We use the first VKITTI sequence (in all weather conditions) as the validation set, and the remaining four sequences as the training set. We consider the tasks of (i) single-frame depth estimation, and (ii) egomotion estimation from pairs of frames.

**Evaluation metrics.** We evaluate the error between our camera motion predictions using four metrics, as defined in prior works [25, 15]. These are: (a) distance error: the camera end point error distance in meters; (b) rotation angle error: the camera rotational error in radians, (c) angular translation error: the error in the angular direction of the camera translation; and (d) magnitude translation error: the error in the magnitude of the camera translation.

We evaluate the error between our depth prediction and the true depth with an L1 distance in log depth space. The use of log depth is inspired by Eigen *et al.* [7], but we do not use the scale-invariant error from that work, because it assumes the presence of an oracle indicating the true mean depth value at test time.

We consider two supervision setups: (a) unsupervised

learning of SfM, where the AIGN is trained from scratch using (self-supervised) reconstruction and adversarial losses, and (b) supervised pretraining, where depth and camera motion in the training set are used to pretrain our generators, before applying unsupervised learning in the test set.

**Self-supervised learning of SfM.** We compare the following models: (a) our full model, *AIGN*, with adversarial losses on depth and camera motion, (b) our model with adversarial losses only on depth, but not on camera motion (*depth-AIGN*), (c) our model without adversarial losses but instead with a smoothness loss on depth (*smooth-IGN*), and (d) our model without any adversarial priors (*IGN*).

We show depth, camera error and photometric errors against number of training iterations for all models in Figure 3 in the test set. Models without depth adversarial priors diverge after a number of iterations (the depth map takes very large values). This is due to the well known *scale ambiguity* of monocular reconstruction: objects can be further away or simply be smaller in size, and the 2D re-projection will be the same. Adversarial priors enforce constraints regarding the general scale of the world scene, as well as depth-like appearance and camera motion-like plausibility on the predictions, and prevent depth divergence. While the adversarial model has higher photometric error than the one without priors (naturally, since it is more constrained), the intermediate variables, namely depth and camera motion, are much more accurate when adversarial priors are present. The model that uses depth smoothness (*depth-IGN*) falls in-between the AIGN model and the model without any priors, and diverges as well. In Figure 4, we show the estimated depth and geometric flow predicted by models with and without adversarial priors.

**Supervised pretraining.** In this setup, we pretrain our camera pose estimation sub-network and depth sub-network using the VKITTI training set using supervised training against ground-truth depthmaps and camera motions supplied with the dataset (*pretrain* baseline). We then finetune our model in the test set using self-supervision, *i.e.*, reprojection error and adversarial constraints (*Pretrain+SelfSup*). We evaluate camera pose estimation performance of our models in the test set, and compare against the geometric model of Jaegle *et al.* [15], a monocular camera motion estimation method that takes optical flow as input, and solves for the camera motion with an optimization-based method that attenuates the influence of outlier flows. We show our translation and rotation camera motion errors in Table 3. The pretrained model performs worse than the geometric baseline. When pretraining is combined with self-supervision, we obtain a much lower error than the geometric model. Monocular geometric methods, such as [15], do not have a good way to resolve the scale ambiguity in reconstruction, and thus have large translation errors. The AIGN for SfM, while being monocular, learns depth and

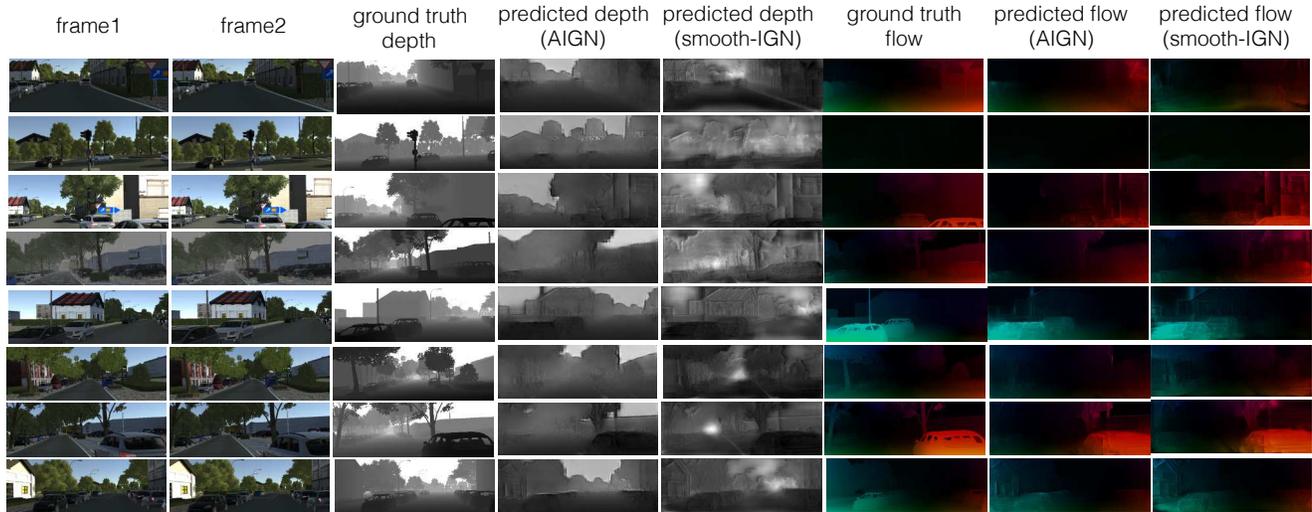


Figure 4. **Structure from motion results with and without adversarial priors.** The results of the baseline (columns 5th and 8th) are obtained from a model with depth smoothness prior (*smooth-IGN*), trained with early stopping at 40K iterations (before divergence).

does not suffer from such ambiguities. Further, we outperform [15] with respect to the angle of translation that a geometric method can in principle estimate (no ambiguity). These results suggest that our adversarial SfM model improves by simply watching unlabeled videos, without diverging.

	trl err	rot err	trl mag.	trl ang.
Geometric [15]	0.4588	<b>0.0014</b>	0.4579	0.3423
Pretrain	0.4876	0.0017	0.4865	<b>0.3306</b>
Pret.+SelfSup	<b>0.1294</b>	<b>0.0014</b>	<b>0.1247</b>	0.3333

Table 3. **Camera motion estimation in our Virtual KITTI test set.** The self-supervised model outperforms the geometric baseline of Jaegle *et al.* [15]. The translation error (column 1) is decomposed into magnitude and angular error in columns 3-4.

### 3.3. Image-to-image translation

We use the CelebA dataset [19] which contains 202,599 face images, with 10,177 unique identities, and is annotated with 40 binary attributes. We preprocess the data by cropping each face image to the largest bounding box that includes the whole face using the OpenFace library [1].

**Biased super-resolution.** We train female-to-male and male-to-female gender transformation by applying adversarial super-resolution to new face input images, while discriminator memories contain only male or only female faces, respectively. We train old-to-young and young-to-old age transformations by applying adversarial super-resolution to new face images while discriminator memories contain only young or only old faces, respectively – as indicated by the attributes in the CelebA dataset. We train identity mixing transformations by applying adversar-

ial super-resolution to new face images while discriminator memories contain only a particular person identity, for demonstration we choose Tom Cruise. We show results in Figure 5 (a-c). We further compare our model against the recent work of Attribute2Image[31] in Figure 5(d) using code available by the authors. The AIGN better preserves the fidelity of the transformation and is more visually detailed. Though we demonstrate age, gender transformation and identity mixing, AIGN could be used for any creative image generation task, with appropriate curation of the discriminator’s ground-truth memories.

**Biased Inpainting.** We train “bigger lips” and “bigger nose” transformations by applying adversarial inpainting to new input face images where the mouth or nose region has been masked, respectively, and discriminator’s memories contain only face images with big lips or with big noses, respectively. Note that “big lips” and “big nose” are attributes annotated in the CelebA dataset. We show results in Figure 6. From top to bottom, we show the original image, the masked image input to adversarial inpainting, the output of our generator, and in the last row, the output of our generator superimposed over the original image.

**Renderers versus parametric decoders for image-to-image translation.** We compare the results of domain-specific, non-parametric renderers, and parametric decoders in image-to-image translation tasks. For our model with parametric decoder we use the full res image as input and instead of the downsampling module in the proposed super-resolution renderer we instead use convolutional/deconvolutional layers so that the decoder can freely adjust its weights through training. This is similar to the one way transformation proposed in the concurrent work of [35]. We trained both models on gender transforma-

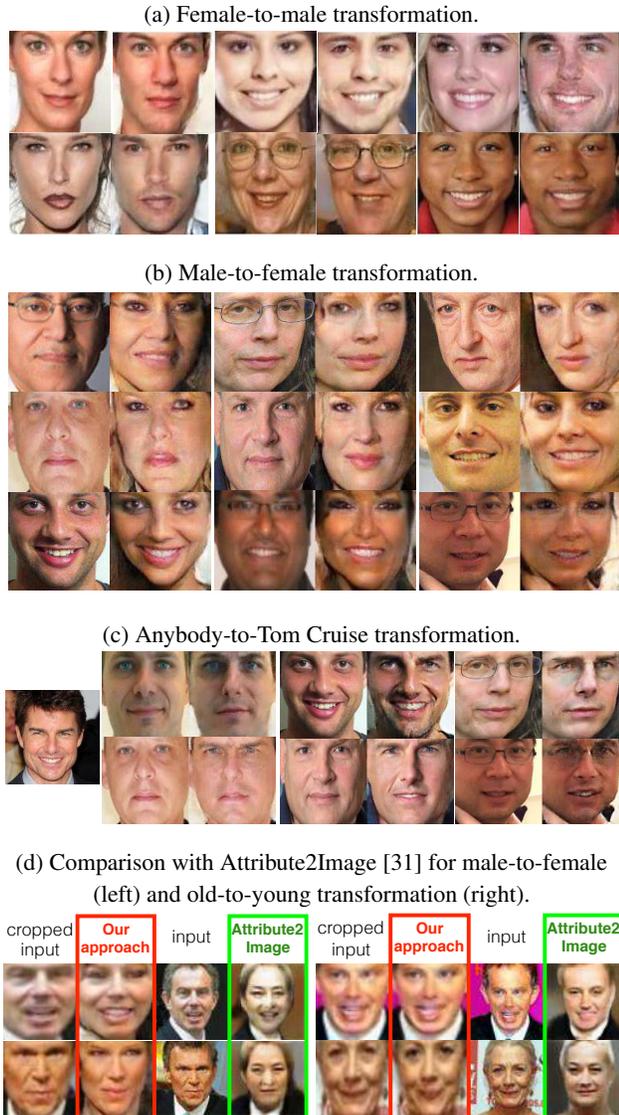


Figure 5. **Biased adversarial super-resolution.**



Figure 6. **Biased adversarial inpainting** for bigger lips transformation (left) and bigger nose transformation (right). From left to right: the original image, the masked image input to our model, and the generated output superimposed over the masked image.

tion, and results are shown in Figure 7. While both models produce photorealistic results, the model with the ren-

derer produces females “more paired” to their male counterparts, while parametric renderers may alter other properties of the face considerably, e.g., in the last row of Figure 7, the age of the produced females does not match their male origins. Parameter-free rendering is an important element of unsupervised learning with AIGNs; We have observed that parametric decoders (instead of parameter-free renderers) can cause the reconstruction loss to drop without learning meaningful predictions but rather exploiting the capacity of the decoder. We provide a comprehensive experiment in the supplementary material.

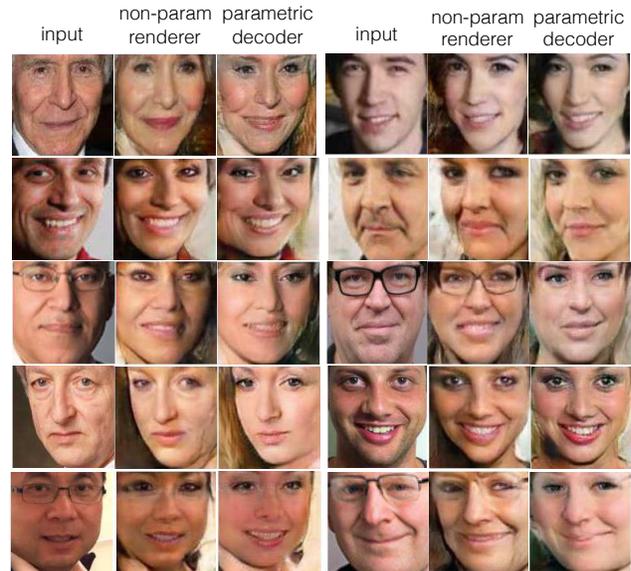


Figure 7. Comparing male-to-female transformation using renderers and parametric decoder.

## 4. Conclusion

We have presented Adversarial Inverse Graphics Networks, weakly supervised neural networks for 2D-to-3D lifting and image-to-image translation that combines feedback from renderings of predictions with data-driven priors on latent semantic factors, imposed using adversarial networks. We showed AIGNs outperform previous supervised models that do not use adversarial priors, in the tasks of 3D human pose estimation and 3D structure and egomotion extraction. We further showed how biasing discriminators’ priors for inpainting and super-resolution results in creative image editing, and outperforms supervised variational autoencoder models of previous works, in terms of the fidelity of the transformation, and the amount of visual detail. Deep neural networks have shown that we do not need to engineer our features; AIGNs shows we do not need to engineer our priors either.

## References

- [1] B. Amos, B. Ludwiczuk, and M. Satyanarayanan. Openface: A general-purpose face recognition library with mobile applications. Technical report, CMU-CS-16-118, CMU School of Computer Science, 2016.
- [2] J. Bruna, P. Sprechmann, and Y. LeCun. Super-resolution with deep convolutional sufficient statistics. *CoRR*, abs/1511.05666, 2015.
- [3] A. Byravan and D. Fox. SE3-Nets: Learning rigid body motion using deep neural networks. *CoRR*, abs/1606.02378, 2016.
- [4] K. Dallenbach. A puzzle-picture with a new principle of concealment. *American Journal of Psychology*, 64(3):431–433, 1951.
- [5] C. Dong, C. C. Loy, K. He, and X. Tang. Image super-resolution using deep convolutional networks. *CoRR*, abs/1501.00092, 2015.
- [6] A. Dosovitskiy, P. Fischer, E. Ilg, P. Häusser, C. Hazirbas, V. Golkov, P. Smagt, D. Cremers, and T. Brox. FlowNet: Learning optical flow with convolutional networks. In *ICCV*, 2015.
- [7] D. Eigen, C. Puhrsch, and R. Fergus. Depth map prediction from a single image using a multi-scale deep network. In *NIPS*, pages 2366–2374, 2014.
- [8] A. Gaidon, Q. Wang, Y. Cabon, and E. Vig. Virtual worlds as proxy for multi-object tracking analysis. In *CVPR*, 2016.
- [9] R. Garg, B. V. Kumar, G. Carneiro, and I. Reid. Unsupervised cnn for single view depth estimation: Geometry to the rescue. In *European Conference on Computer Vision*, pages 740–756. Springer, 2016.
- [10] A. Geiger, P. Lenz, and R. Urtasun. Are we ready for autonomous driving? The KITTI vision benchmark suite. In *CVPR*, 2012.
- [11] A. Handa, M. Bloesch, V. Pătrăucean, S. Stent, J. McCormac, and A. Davison. gynn: Neural network library for geometric computer vision. In *Computer Vision—ECCV 2016 Workshops*, pages 67–82. Springer International Publishing, 2016.
- [12] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, 2016.
- [13] C. Ionescu, D. Papava, V. Olaru, and C. Sminchisescu. Human3.6m: Large scale datasets and predictive methods for 3d human sensing in natural environments. *TPAMI*, 36(7), 2014.
- [14] P. Isola, J. Zhu, T. Zhou, and A. A. Efros. Image-to-image translation with conditional adversarial networks. *CoRR*, abs/1611.07004, 2016.
- [15] A. Jaegle, S. Phillips, and K. Daniilidis. Fast, robust, continuous monocular egomotion computation. In *ICRA*, pages 773–780, 2016.
- [16] C. Kerl, J. Sturm, and D. Cremers. Dense visual SLAM for RGB-D cameras. In *IROS*, 2013.
- [17] T. D. Kulkarni, W. Whitney, P. Kohli, and J. B. Tenenbaum. Deep convolutional inverse graphics network. *CoRR*, abs/1503.03167, 2015.
- [18] C. Ledig, L. Theis, F. Huszar, J. Caballero, A. P. Aitken, A. Tejani, J. Totz, Z. Wang, and W. Shi. Photo-realistic single image super-resolution using a generative adversarial network. *CoRR*, abs/1609.04802, 2016.
- [19] Z. Liu, P. Luo, X. Wang, and X. Tang. Deep learning face attributes in the wild. In *Proceedings of International Conference on Computer Vision (ICCV)*, 2015.
- [20] D. Pathak, P. Krähenbühl, J. Donahue, T. Darrell, and A. Efros. Context encoders: Feature learning by inpainting. In *CVPR*, 2016.
- [21] V. S., R. S., S. C., S. R., and F. K. Sfm-net: Learning of structure and motion from video. In *arxiv*, 2017.
- [22] T. Schöps, J. Engel, and D. Cremers. Semi-dense visual odometry for AR on a smartphone. In *ISMAR*, 2014.
- [23] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- [24] C. K. Sønderby, J. Caballero, L. Theis, W. Shi, and F. Huszár. Amortised map inference for image super-resolution. *arXiv preprint arXiv:1610.04490*, 2016.
- [25] J. Sturm, N. Engelhard, F. Endres, W. Burgard, and D. Cremers. A benchmark for the evaluation of RGBD SLAM systems. In *IROS*, 2012.
- [26] F. Tung and K. Fragkiadaki. Inverse problems in computer vision with adversarial imagination priors. *openreview*, 2016.
- [27] S.-E. Wei, V. Ramakrishna, T. Kanade, and Y. Sheikh. Convolutional pose machines. In *CVPR*, 2016.
- [28] J. Wu, T. Xue, J. J. Lim, Y. Tian, J. B. Tenenbaum, A. Torralba, and W. T. Freeman. Single image 3D interpreter network. In *European Conference on Computer Vision*, pages 365–382, 2016.
- [29] J. Wu, T. Xue, J. J. Lim, Y. Tian, J. B. Tenenbaum, A. Torralba, and W. T. Freeman. Single image 3D interpreter network. In *ECCV*, 2016.
- [30] J. Wu, C. Zhang, T. Xue, W. T. Freeman, and J. B. Tenenbaum. Learning a probabilistic latent space of object shapes via 3d generative-adversarial modeling. In *NIPS*, 2016.
- [31] X. Yan, J. Yang, K. Sohn, and H. Lee. Attribute2image: Conditional image generation from visual attributes. *CoRR*, abs/1512.00570, 2015.
- [32] R. Yeh, C. Chen, T. Lim, M. Hasegawa-Johnson, and M. N. Do. Semantic image inpainting with perceptual and contextual losses. *arXiv preprint arXiv:1607.07539*, 2016.
- [33] T. Zhou, M. Brown, N. Snavely, and D. G. Lowe. Unsupervised learning of depth and ego-motion from video. In *arxiv*, 2017.
- [34] X. Zhou, M. Zhu, G. Pavlakos, S. Leonardos, K. G. Derpanis, and K. Daniilidis. Monocap: Monocular human motion capture using a CNN coupled with a geometric prior. *CoRR*, abs/1701.02354, 2017.
- [35] J.-Y. Zhu, T. Park, P. Isola, and A. A. Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. *arXiv preprint arXiv:1703.10593*, 2017.