

# Hallucinating Pose-Compatible Scenes

Tim Brooks      Alexei A. Efros

UC Berkeley

**Abstract.** What does human pose tell us about a scene? We propose a task to answer this question: given human pose as input, hallucinate a compatible scene. Subtle cues captured by human pose — action semantics, environment affordances, object interactions — provide surprising insight into which scenes are compatible. We present a large-scale generative adversarial network for pose-conditioned scene generation. We significantly scale the size and complexity of training data, curating a massive meta-dataset containing over 19 million frames of humans in everyday environments. We double the capacity of our model with respect to StyleGAN2 to handle such complex data, and design a pose conditioning mechanism that drives our model to learn the nuanced relationship between pose and scene. We leverage our trained model for various applications: hallucinating pose-compatible scene(s) with or without humans, visualizing incompatible scenes and poses, placing a person from one generated image into another scene, and animating pose. Our model produces diverse samples and outperforms pose-conditioned StyleGAN2 and Pix2Pix/Pix2PixHD baselines in terms of accurate human placement (percent of correct keypoints) and quality (Fréchet inception distance).

## 1 Introduction

Human pose can reveal a lot about a scene. For example, mime artists<sup>1</sup> invoke vivid scenes in a viewer’s mind through pose and movement alone, despite performing on a bare stage. The viewer is able to imagine the invisible objects and scene elements because of the strong relationship between human poses and scenes learned through a lifetime of daily observations.

Psychologists have long been interested in understanding this symbiotic relationship between human and scene [24,6]. J.J. Gibson proposed the notion of *affordances* [24], which can be described as “opportunities for interactions” furnished by the environment. In computer vision, affordances have been used to provide a functional description of the scene. Given an image, a number of approaches try to predict likely human poses these scenes afford [29,15,23,30].

This work, on the other hand, considers the opposite problem: given a human pose as input, the goal is to hallucinate scene(s) that are compatible with that pose. Consider Figure 1. A push-up pose (top) places severe constraints on the space of compatible scenes: they must not only be semantically compatible (e.g.,

---

<sup>1</sup> For those unfamiliar with mime artists, here is a wonderful example performance: [https://youtu.be/FPMBV3rd\\_hI](https://youtu.be/FPMBV3rd_hI)

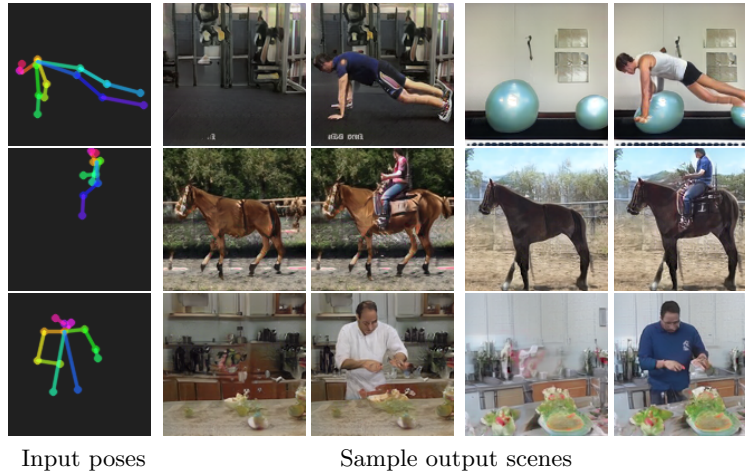


Fig. 1: Given a human pose as input, the goal of this paper is to hallucinate scene(s) that are compatible with that pose. Our model can generate isolated scenes as well as scenes containing humans.

gym, exercise room), but also have compatible spatial affordances (enough floor space or appropriate equipment). Objects in the scene can afford interaction with the human (e.g., squishing down an exercise ball). Other poses might not appear as constraining, but even a simple standing pose (bottom) — head looking down, hands reaching in, legs occluded — is actually a strong indicator of a cooking scene, and signals that an object (e.g., countertop) must be occluding the legs.

Rather than explicitly model scene affordances and contextual compatibility, we employ a modern large-scale generative model (based on a souped-up StyleGAN2 [44] architecture) to *discover* these relationships implicitly, from data. While GANs have performed well at capturing disentangled visual models in specialized scenarios (e.g., faces, churches, categories from ImageNet [16]), they have not been demonstrated *in situ*, on complex, real-world data across varying environments.

We curate a massive meta-dataset of humans interacting with everyday environments, containing over 19 million frames. The complexity and scale of data is much higher than common GAN datasets, such as FFHQ [43] (70,000 face images) and ImageNet [16] (1.3M object images). With an appropriate pose conditioning mechanism, increased model capacity, and removal of style mixing, we are able to successfully train a pose-conditioned GAN on this highly complex data. Our model and meta-dataset mark substantial progress leveraging GANs in real-world settings containing humans and diverse environments. Through numerous visual experiments, we demonstrate our model’s emergent ability to capture affordances and contextual relationships between poses and scenes.

See our webpage<sup>2</sup> for our supplemental video and code release.

<sup>2</sup> <https://www.timothybrooks.com/tech/hallucinating-scenes>

## 2 Related Work

*Scene and object affordances.* Affordances [24] describe the possible uses of a given object or environment. A significant body of work learns scene affordances, such as where a person can stand or sit, from observing data of humans [29,28,21,15,38,23,74,14,49]. Overlapping areas of work focus on human interactions with objects [77,47,82,25,10] or synthesize human pose conditioned on an input scene [48,8,71]. We propose the reverse task of hallucinating a scene conditioned on pose.

*Pose-conditioned human synthesis.* There are a plethora of methods that take a source image (or video) of a human plus a new pose and generate an image of the human in the new pose [53,67,4,72,1,13,50]. Although we too condition on pose, our goals are almost entirely opposite: we aim to generate novel scenes compatible with a given pose, whereas the above methods reuse the scene from the source image/video and only focus on reposing within that provided scene.

*GANs for image synthesis.* Introduced by Goodfellow *et al.* [27] a generative adversarial network (GAN) is an implicit generative model that learns to synthesize data samples by optimizing a minimax objective. The generator is tasked with fooling a discriminator, and the discriminator is tasked with differentiating real and generated samples. Modern GANs are capable of producing high quality images [7,41,43,44]. Image translation [36,72] utilizes conditional GANs [58] to translate from one domain to another. While our task is pose-conditional scene generation, we leverage benefits of modern unconditional GANs [44].

*Visual disentanglement.* Disentanglement methods attempt to separate out independent controllable attributes of images. This can be achieved with unsupervised methods [43,34,37,63], or an auxiliary signal [26,51]. Components of image samples can be added, removed and composed using pretrained GANs [5,12]. Recent work has applied similar strategies to image translation models to compose style and content from different images [62]. The most related to us is the work of Ma *et al.* [54], who synthesize images of people, while independently controlling foreground, background, and pose. However, the focus is on generating humans in very tightly cropped images with simple backgrounds, rather than generating scenes with appropriate affordances. Many disentanglement methods assume all images or image attributes can be combined with all others [43,34,26,37,63,62,51,54]. In this work, we seek disentangled representations of pose, human appearance and scene, yet it is essential our model understand which scenes can or cannot be composed with which poses.

*Contextual relationships.* Many works leverage contextual relationships among objects and scenes [6] to improve vision models such as object recognition and semantic segmentation [70,64,18,61]. Divvala *et al.* [18] explicitly enumerate (Table 1) a taxonomy of possible contextual information. In this paper we are specifically interested in contextual relationships between humans and their environments, and aim to recover them implicitly, from data.



Fig. 2: **Dataset comparison.** (a) The largest human-centric video dataset with ground truth poses uses a fixed background, missing scene interactions. (b) Action recognition datasets include scenes, but contain videos without people or of close-up content. (c) Our dataset is a massive curation of humans in scenes.

### 3 Humans in Context Meta-dataset

To study the rich relationship between scenes and human poses requires large-scale data of people interacting with many different environments. Internet videos are a natural source, containing vast data of daily human activities. Unfortunately, large-scale action recognition datasets [17,60,45] include substantial content without humans, as well as close-up footage not of scenes. Most existing human-centric datasets are insufficiently small [68,3], narrow in scene type [22,80], or captured on a fixed background [35].

We therefore curate a meta-dataset of 229,595 video clips, each containing a single person in a scene, sourced from 10 existing human and action recognition video datasets [17,60,45,68,40,19,3,22,80,76], and supplemented with pseudo-ground truth pose obtained using OpenPose [11,9]. Video offers a massive source of real-world data, and ensures all poses of human activity are represented, rather than only poses photographers choose to capture in still images.

Videos are extensively filtered for quality, ensuring satisfactory framerate, bitrate and resolution. 1,509,032 videos (75% of source videos) pass quality filtering. Frames are then filtered with pretrained Keypoint R-CNN [31,75] person detection and OpenPose [11,9] keypoint prediction models. The final dataset only includes clips of at least 30 frames where Keypoint R-CNN detects a single person and OpenPose predicts sufficient keypoints. This results in 19,503,700 frames (7.8% of high quality frames), with each clip averaging 85 frames long.

While we train on images, we split data into partitions based on video clips, reserving 12,800 clips for testing and the remaining 216,795 for training. See the supplement for dataset details.

### 4 Pose-compatible Scene GAN

We design a conditional GAN [27,58] to produce scenes compatible with human pose. Our network architectures are based on StyleGAN2 [44] and are depicted in Figure 3. Generating high quality pose-compatible scenes arises from simple yet important modifications: dual pose conditioning, removal of style mixing, and large-scale training. Our model can produce isolated scene images without any human by zeroing out keypoint heatmaps when generating images.






Fig. 3: **Network architectures.** Our networks are based on StyleGAN2, with simple modifications to ensure accurately placed humans and compatible scenes. In particular, the conditional generator and discriminator networks utilize pose  $p$  via two mechanisms: keypoint heatmaps and pose latent conditioning. Keypoint heatmaps correctly position a human, and pose latent conditioning drives latent codes  $w$  to generate compatible scenes. Multiple plausible scenes can be produced for the same input pose by sampling different noise vectors  $z$ .

#### 4.1 Dual pose conditioning

The conditional generator  $G$  and discriminator  $D$  both utilize input pose via two mechanisms: keypoint heatmap conditioning, which specifies spatial placement of a human subject, and pose latent conditioning, which infers compatible scenes. To succeed at our task, humans must be positioned correctly and generated scenes must be compatible. Dual pose conditioning drives strong performance in both respects, and outperforms conditioning on either alone in our ablation experiment (Table 3). Furthermore, dual pose conditioning disentangles control of scene and human pose. We leverage these separate controls for numerous applications: generating scenes without humans, visualizing incompatible scenes and poses, placing a person in a new scene, and animating pose.

*Keypoint heatmaps.* Let pose  $p = (p_1, \dots, p_K)$  denote 2D locations of the  $K = 18$  human keypoints detected by OpenPose [9], and let  $v = (v_1, \dots, v_K)$  indicate visibility of each keypoint. Following the works of [67, 4, 1], our keypoint heatmaps  $H(p)$  consist of radial basis function kernels centered at each keypoint. For heatmap  $k \in \{1, \dots, K\}$ , the intensity at location  $q$  is given by Equation 1. We concatenate heatmaps at each scale of the generator, and at the input of the discriminator. We set  $\sigma^2 = \max(0.5, 0.005R^2)$  where  $R$  is the spatial resolution of the heatmaps. After training, we generate images of scenes without humans by simply zeroing out all keypoint heatmaps.

$$H_{k,q}(p) = \begin{cases} \exp\left(-\frac{\|q - p_k\|^2}{2\sigma^2}\right) & \text{if } v_k = 1 \\ 0 & \text{otherwise} \end{cases} \quad (1)$$

*Pose latent conditioning.* To generate compatible scenes, we condition scene latent codes on the input pose. Akin to intermediate latents in StyleGAN2 [44], the scene latent code  $w$  controls generation by modulating convolutional weights.

To condition the latent code, pose locations and visibility are flattened and mapped to a 512-dimensional input via a learned linear projection. A noise sample  $z \sim \mathcal{Z}$  is concatenated with the input vector and passed through a multi-layer perceptron (MLP)  $f_G$  to produce a scene latent code  $w \in \mathcal{W}$ . Multiple plausible scenes can be generated by sampling different noise vectors  $z$  for the same pose. The discriminator learns a separate linear projection and MLP  $f_D$ .

## 4.2 Removal of style mixing

Style mixing regularization [43,62] encourages disentanglement by randomly mixing intermediate latent codes during training. The technique assumes image attributes at each layer are compatible with all other image attributes (e.g. any face could have any color hair). This assumption is not true when composing scenes and humans, which we visually demonstrate through the incompatible scenes and poses in Figure 6. This motivates removing style mixing regularization during training, which improves results in our ablation experiments (Table 2).

## 4.3 Large-scale GAN training

Typical datasets used with StyleGAN2 (e.g. faces, bedrooms, churches [52,78]) are relatively homogeneous. Increasing model capacity is a natural extension given the diversity and complexity of scene images in our dataset. We find that increasing the channel width of convolutional layers by  $2\times$  significantly improves our model (see ablation in Table 2). Following prior work in scaling GANs [7], we also increase minibatch size (from 40 to 120). Concurrent work [59,66] also explores scaling StyleGAN, and proposes strategies such as self-filtering the training dataset [59], progressive growing and leveraging pretrained classifiers [66].

## 4.4 Model details

We train all models at  $128 \times 128$  resolution with non-saturating logistic loss [27], path length [44] and  $R_1$  [57] regularization, and exponential moving average of generator parameters [41]. We remove spatial noise maps from StyleGAN2 for all models, employ differentiable augmentation of both real and generated images [81,42], and train the discriminator with an additional fake example containing real images with mismatched labels [65]. See the supplement for details.

# 5 Experiments

Our model hallucinates diverse, high quality images of scenes compatible with input pose. We generate scenes in isolation as well as scenes containing humans, and analyze our model through several visual experiments. Generating scene images is challenging due to the high complexity of data, and our model outperforms Pix2Pix/Pix2PixHD [36,73] and pose-conditioned StyleGAN2 [44] baselines in terms of image quality and accurate human placement. We present characteristic success and failure results in Figure 4 and Figure 5. See the supplement for more results, including multiple pages of random uncured samples.

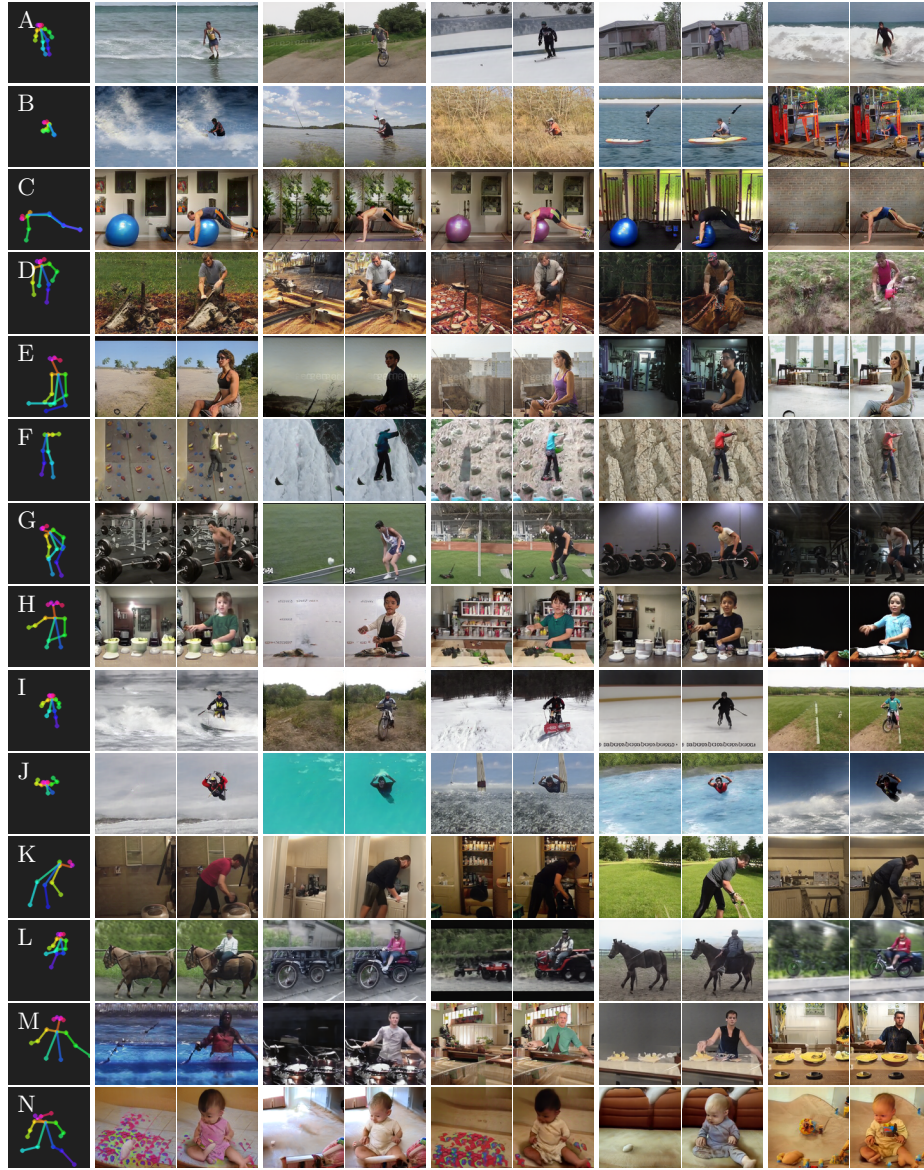


Fig. 4: **Success cases.** Our model learns complex scene-pose relationships. For each input pose, we show many hallucinated scenes, with and without a human. Diverse outputs include a person paddling a kayak (B), lifting a barbell in their hand (G), cleaning the toilet (K), and playing the drums (M). Our model produces multiple plausible scenes for the same pose, providing insight into scenes with related affordances: in the same pose, a person may climb in an indoor gym or on a snowy ledge (F); a person can ride a horse, ride a bicycle, or ride a tractor (L). Please see the appendix for multiple pages of random results.



Fig. 5: **Failure cases.** Causes for failure include: partially generating objects, such as a bike (A); poor overall image quality (B); missing limbs without proper occluders (C); difficulty placing objects, such as a golf club, in a person’s hands (D); difficulty hallucinating an object on which to sit (E); overly repetitive textures (F); infeasible scenes, such as walking on water (G); and leaving behind a partial human when hallucinating the scene in isolation (H).

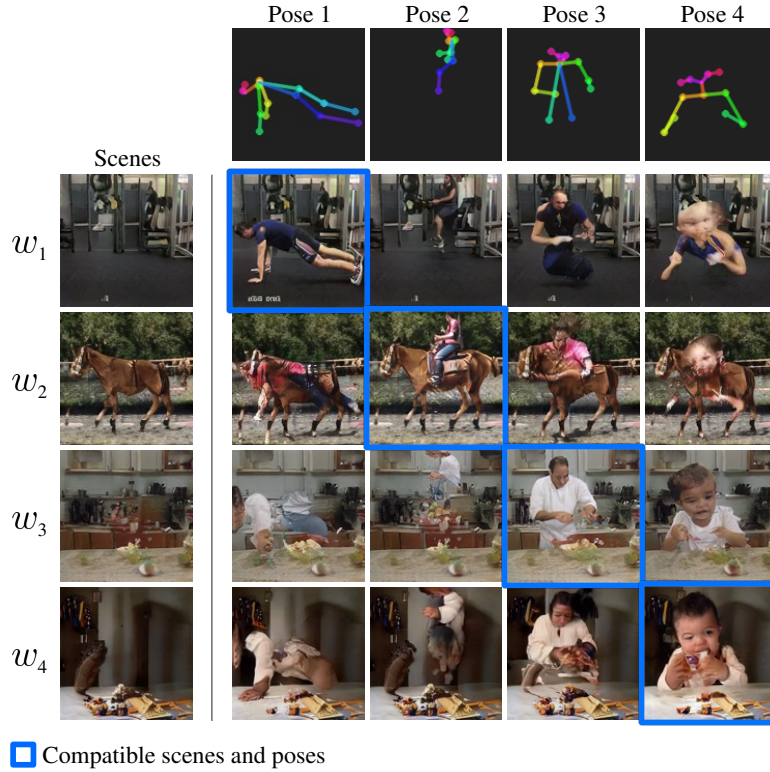


Fig. 6: A central theme of our paper is that scenes must be compatible with human poses to produce realistic images — here we visualize what happens when scenes and poses are *not* compatible. Correctly paired images are shown in blue on the diagonal — a person doing a pushup in a gym, riding a horse, cooking in a kitchen, and a baby leaning on a table. These exemplify interesting relationships between human pose and scene learned by our model. Other images mix scene latent codes with keypoint heatmaps from the wrong pose, often producing unrealistic images. Generating pose-compatible scenes is essential to avoid these incorrect pairings.




### 5.1 Not all scenes and poses are compatible

It is essential that we model which scenes are compatible with which poses. A person cannot do a push-up in the middle of a horse, ride atop a kitchen countertop, or be occluded by thin air. These scenarios sound obviously false, yet could occur if the scene and human pose are incompatible. We visualize images generated with correctly and incorrectly paired scenes and poses in Figure 6.


These examples of incompatible scenes and poses highlight an important difference between our scene data and other datasets commonly used for GAN training, such as cropped faces in the CelebA [52] and FFHQ [43] datasets. Any face can be given glasses, longer or shorter hair, or a darker or lighter skin tone and still remain a feasible image. This enables global disentanglement of attributes, and applications like style mixing, which combines different intermediate latent codes of any two samples (see Figure 3 of the original StyleGAN paper [43] for a wonderful example). The assumption of compatibility between all attribute pairs no longer holds for data of scenes with humans, which motivates conditioning scene latent codes on pose. Relatedly, we find that removing style mixing from training significantly improves performance (Table 2).

### 5.2 Scene occlusion reasoning

Portions of a human pose may be occluded by foreground objects, such as pieces of furniture. Provided a partially visible human pose, our model hallucinates scenes with foreground objects to occlude portions of the pose not visible. Figure 7a shows an example full-body pose and output scenes. When the legs are not visible in the input pose in 7b, our model produces scenes with occluders blocking the legs, demonstrating its emergent ability to reason about occlusions.



(a) Scenes generated with a full-body input pose.



(b) Scenes generated from the same input pose with legs not visible.

Fig. 7: (a) A full-body input pose and corresponding scenes. (b) When the legs from an otherwise identical pose are hidden, our model hallucinates scenes with foreground objects, such as a drum kit or table, to occlude the missing legs.




Fig. 8: Given an input pose (top left), our method can compose human appearances (top row) and scenes (left column) from different generated images.



Fig. 9: Provided an input pose sequence (top), we infer scenes based on the first pose, then generate animations (middle/bottom) by keeping the scene latent fixed and passing keypoint heatmaps for each subsequent pose.

### 5.3 Human appearance and scene disentanglement

Section 5.1 demonstrates why complete separation of pose and scene is undesirable. We can, however, disentangle human appearance from scene when both are conditioned on the same pose, as shown in Figure 8. To achieve this, we optimize a latent code to compose two samples. We minimize perceptual loss [39,79] between person-only crops of the composition and first sample, and scene-only images of the composition and second sample. See the supplement for details.

### 5.4 Animating pose

After training, our model is capable of animating pose in a stationary scene. In Figure 9 we demonstrate a sequence of images generated by fixing the scene and animating the human pose. The scene is inferred from only the first pose, and is limited to small human motion and stationary backgrounds.



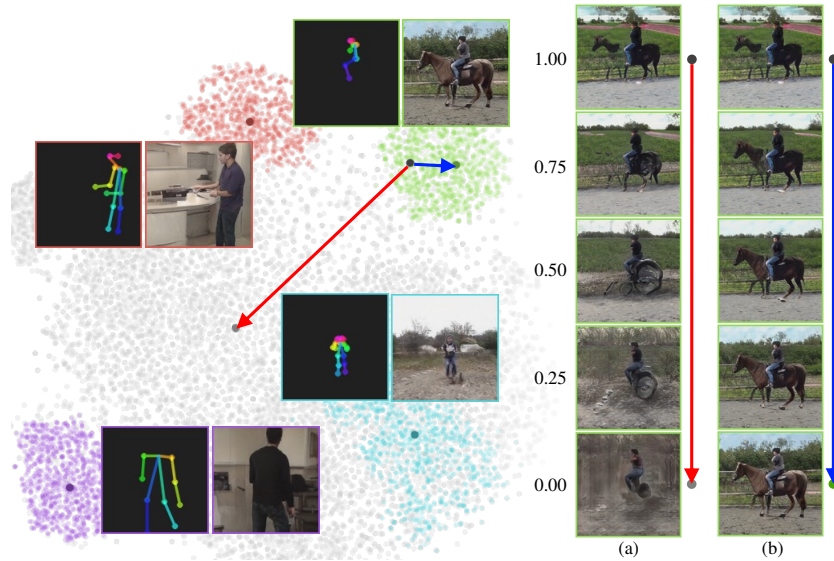


Fig. 10: We contrast truncation via (a) interpolation toward the mean of random latents, and (b) interpolation toward the mean of conditional latent clusters. The left plot shows a t-SNE [55] visualization of latent codes. Gray points are 10,000 random latents. Colored sets of points are each 1000 latent samples conditioned on the same pose. The formation of clusters signifies that different scene latents conditioned on the same pose are close to each other in the intermediate latent space. The dark gray point in the center is the mean of all random latents, and dark colored points are the means for each pose. Beside each cluster is the input pose and image generated using the corresponding mean cluster latent. Conditional truncation (b) works significantly better than unconditional (a).

### 5.5 Scene clustering and truncation

Regions of low density in the data distribution are particularly challenging to model. Quality can be improved (at the loss of some diversity) by sampling from a shrunk distribution [2, 56, 46, 7, 20, 33]. StyleGAN [43] interpolates intermediate latents  $w$  toward the mean  $\bar{w} = \mathbb{E}_{z \sim \mathcal{Z}}[w]$  to shrink the sampling distribution, which improves generation quality for models trained on data such as faces. However, on our more complex data, interpolating toward the mean scene latent produces a gray scene rather than improving quality, as shown in Figure 10a.

In visualizing a t-SNE [55] plot of scene latents in Figure 10, we observe that latents sampled from different noise vectors  $z$  yet conditioned on the same pose  $p$  form clusters. We apply conditional truncation by interpolating a latent  $w$  toward the conditional mean  $\bar{w}_p = \mathbb{E}_{z \sim \mathcal{Z}}[w|p]$ , shifting the sample toward the cluster center. Shown in part (b) of Figure 10, conditional truncation works significantly better for our model. We apply conditional truncation  $w' = \bar{w}_p + \psi(w - \bar{w}_p)$  of  $\psi = 0.75$  to generated images throughout the paper. Concurrent work [59] proposes a similar method for applying truncation toward the centers of perceptual clusters.



Fig. 11: Baseline comparisons. Pix2Pix/Pix2PixHD struggle to produce realistic images; pose-conditioned StyleGAN2 (SG2) often generates humans in the wrong pose; our model generates realistic scenes with humans in the correct pose.

## 5.6 Baseline comparisons

Please see Figure 11 for visual comparisons with baseline methods. Pix2Pix and Pix2PixHD were designed for image translation tasks with stronger conditioning, such as segmentation masks. These methods struggle to produce reasonable images on our more challenging task and dataset. StyleGAN2 (SG2) with latent pose conditioning provides a stronger baseline, but still has notable issues with image quality and often places humans in the incorrect pose. These observations are corroborated by metric performance in two respects: how accurately human subjects are positioned, and how realistic generated scenes look.

To succeed at our task, a model must both put a human in the correct pose and generate a compatible scene. Table 1 compares our model with Pix2Pix, Pix2PixHD and StyleGAN2 baselines on these metrics, demonstrating that our model achieves superior performance. Note that StyleGAN2 [44] is primarily an unconditional GAN. The public code release and follow-up work [42] support class-conditional generation. We refer to the version of our model with only pose latent conditioning as StyleGAN2, since it is the most straightforward extension of StyleGAN2 for our task.

*Accurate human positioning.* PCKh [3] measures the percent of correct pose keypoints (within a radius relative to the head size), where a higher percent is better. We use OpenPose [9] to extract poses from generated images for comparison with input poses. PCKh is computed on a held out test set, ensuring accurate placement of new poses not seen during training.

*Realistic scene images.* FID — Fréchet inception distance [32] — measures realism by comparing distributions of Inception network [69] features between the training dataset and generated images. Lower FID scores are better and correlate with higher quality, more realistic images.

Table 1: **Baseline metric comparisons.** We report PCKh (higher is better) as a measure of how accurately humans are positioned, and FID (lower is better) as a measure of how realistic generated scenes look. Our model outperforms Pix2Pix, Pix2PixHD and pose-conditioned StyleGAN2 baselines on both metrics. While the poor performance of baselines may appear surprising, note that our task is much more challenging than standard conditional generation tasks: the dataset is diverse and complex, and conditioning on pose requires the network to infer scene contents and layout.

	PCKh $\uparrow$	FID $\downarrow$
Pix2Pix	48.4	71.2
Pix2PixHD	73.8	149.7
StyleGAN2 (with pose latent conditioning)	32.4	16.6
Ours	84.2	5.9

Table 2: **StyleGAN2 ablation.** We enumerate modifications relative to a pose-conditioned StyleGAN2 baseline. In particular, removing style mixing, conditioning on keypoint heatmaps, augmenting discriminator inputs and passing a fake mismatched example to the discriminator, and increasing scale all contribute to our final model.

	PCKh $\uparrow$	FID $\downarrow$
StyleGAN2 (with pose latent conditioning)	32.4	16.6
– style mixing	36.4	11.6
+ keypoint heatmaps	79.8	12.2
+ augmentation, mismatch	80.7	12.1
+ large scale (Ours)	84.2	5.9

Table 3: **Pose conditioning ablation.** We contrast three options for pose conditioning: only conditioning the latent on pose, only conditioning on keypoint heatmaps, and dual conditioning of both latents and heatmaps. We conduct this ablation on the smaller version of our model. We find that keypoint heatmap conditioning is crucial for accurately placing a human (PCKh), whereas latent conditioning improves the quality of scene generation (FID). We condition with both mechanisms in our final model, which has the best metric trade-off, and enables separating control of human position and scene generation after training.

Conditioning method	PCKh $\uparrow$	FID $\downarrow$
Latent only	36.4	11.6
Heatmap only	79.7	15.1
Both	79.8	12.2

## 5.7 Ablations

We present two ablation experiments. Table 2 enumerates changes relative to a pose-conditioned StyleGAN2 baseline, demonstrating improvements gained by our simple yet important modifications. Table 3 compares three options for pose conditioning: latents only, keypoint heatmaps only, and dual conditioning of both. Keypoint heatmaps are necessary to accurately position a human in the scene, which is shown by a substantially higher PCKh. Latent conditioning improves quality, which is shown by a lower FID score. We condition with both mechanisms — in addition to offering the best trade-off in metric performance, dual conditioning enables applications of disentanglement, such as generating scenes without humans or visualizing incompatible scenes and poses.

## 6 Discussion

*Limitations.* Our dataset and model only consider images with a single human subject. Dataset curation is limited by the performance of Keypoint R-CNN [31,75] and OpenPose [11,9] when filtering videos for humans. Training depends on OpenPose to correctly predict poses. Our model does not consider human movement when inferring scenes.

*Societal impact.* There is some risk of this or future generative models being used to create fake and misleading content. Our model also inherits any demographic bias present in the existing datasets used to source our training data.

*Conclusion.* In this paper, we present a new task: provided a human pose as input, hallucinate the possible scene(s) which are compatible with that input pose. Strong relationships between humans, objects and environments dictate which scenes afford a given pose. Many prior works study human affordances from the angle of predicting possible poses given an input scene — we study the other side of the same coin, and hallucinate scenes that afford an input pose.

We demonstrate the emergent ability of our model to capture affordance relationships between scenes and poses. This work marks a significant step toward using GANs to represent complex real-world environments. We hope it will motivate the broader research community to leverage modern generative approaches for scene understanding and modeling.

*Acknowledgements.* We thank William Peebles, Ilija Radosavovic, Matthew Tan-cik, Allan Jabri, Dave Epstein, Lucy Chai, Toru Lin, Shiry Ginosar, Angjoo Kanazawa, Vickie Ye, Karttikeya Mangalam, and Taesung Park for insightful discussion and feedback. Tim Brooks is supported by the National Science Foundation Graduate Research Fellowship under Grant No. 2020306087. Additional support for this project is provided by DARPA MCS, and SAP.

## References

1. Aberman, K., Shi, M., Liao, J., Lischinski, D., Chen, B., Cohen-Or, D.: Deep video-based performance cloning. In: Computer Graphics Forum. vol. 38, pp. 219–233. Wiley Online Library (2019) [3](#), [5](#)
2. Ackley, D.H., Hinton, G.E., Sejnowski, T.J.: A learning algorithm for boltzmann machines. *Cognitive science* **9**(1), 147–169 (1985) [11](#)
3. Andriluka, M., Pishchulin, L., Gehler, P., Schiele, B.: 2d human pose estimation: New benchmark and state of the art analysis. In: Proceedings of the IEEE Conference on computer Vision and Pattern Recognition. pp. 3686–3693 (2014) [4](#), [12](#)
4. Balakrishnan, G., Zhao, A., Dalca, A.V., Durand, F., Gutttag, J.: Synthesizing images of humans in unseen poses. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 8340–8348 (2018) [3](#), [5](#)
5. Bau, D., Zhu, J.Y., Strobel, H., Zhou, B., Tenenbaum, J.B., Freeman, W.T., Torralba, A.: Gan dissection: Visualizing and understanding generative adversarial networks. In: Proceedings of the International Conference on Learning Representations (ICLR) (2019) [3](#)
6. Biederman, I.: On the semantics of a glance at a scene. In: Perceptual Organization (1981) [1](#), [3](#)
7. Brock, A., Donahue, J., Simonyan, K.: Large scale GAN training for high fidelity natural image synthesis. In: International Conference on Learning Representations (2019), <https://openreview.net/forum?id=B1xsqj09Fm> [3](#), [6](#), [11](#)
8. Cao, Z., Gao, H., Mangalam, K., Cai, Q., Vo, M., Malik, J.: Long-term human motion prediction with scene context. In: ECCV (2020) [3](#)
9. Cao, Z., Hidalgo, G., Simon, T., Wei, S.E., Sheikh, Y.: Openpose: realtime multi-person 2d pose estimation using part affinity fields. *IEEE transactions on pattern analysis and machine intelligence* **43**(1), 172–186 (2019) [4](#), [5](#), [12](#), [14](#)
10. Cao, Z., Radosavovic, I., Kanazawa, A., Malik, J.: Reconstructing hand-object interactions in the wild. arXiv e-prints pp. arXiv–2012 (2020) [3](#)
11. Cao, Z., Simon, T., Wei, S.E., Sheikh, Y.: Realtime multi-person 2d pose estimation using part affinity fields. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 7291–7299 (2017) [4](#), [14](#)
12. Chai, L., Wulff, J., Isola, P.: Using latent space regression to analyze and leverage compositionality in gans. In: International Conference on Learning Representations (2021) [3](#)
13. Chan, C., Ginosar, S., Zhou, T., Efros, A.A.: Everybody dance now. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 5933–5942 (2019) [3](#)
14. Chuang, C.Y., Li, J., Torralba, A., Fidler, S.: Learning to act properly: Predicting and explaining affordances from images. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 975–983 (2018) [3](#)
15. Delaitre, V., Fouhey, D., Laptev, I., Sivic, J., Gupta, A., Efros, A.: Scene semantics from long-term observation of people. In: Proc. 12th European Conference on Computer Vision (2012) [1](#), [3](#)
16. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: ImageNet: A Large-Scale Hierarchical Image Database. In: CVPR09 (2009) [2](#)
17. Diba, A., Fayyaz, M., Sharma, V., Paluri, M., Gall, J., Stiefelhagen, R., Van Gool, L.: Large scale holistic video understanding. In: European Conference on Computer Vision. pp. 593–610. Springer (2020) [4](#)

18. Divvala, S.K., Hoiem, D., Hays, J.H., Efros, A.A., Hebert, M.: An empirical study of context in object detection. In: 2009 IEEE Conference on computer vision and Pattern Recognition. pp. 1271–1278. IEEE (2009) [3](#)
19. Epstein, D., Chen, B., Vondrick, C.: Oops! predicting unintentional action in video. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 919–929 (2020) [4](#)
20. Fan, A., Lewis, M., Dauphin, Y.: Hierarchical neural story generation. arXiv preprint arXiv:1805.04833 (2018) [11](#)
21. Fouhey, D.F., Delaitre, V., Gupta, A., Efros, A.A., Laptev, I., Sivic, J.: People watching: Human actions as a cue for single view geometry. In: European Conference on Computer Vision. pp. 732–745. Springer (2012) [3](#)
22. Fouhey, D.F., Kuo, W.c., Efros, A.A., Malik, J.: From lifestyle vlogs to everyday interactions. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 4991–5000 (2018) [4](#)
23. Fouhey, D.F., Wang, X., Gupta, A.: In defense of the direct perception of affordances (2015) [1](#), [3](#)
24. Gibson, J.J.: The Ecological Approach to Visual Perception. Houghton Mifflin, Boston (1979) [1](#), [3](#)
25. Gkioxari, G., Girshick, R., Dollár, P., He, K.: Detecting and recognizing human-object interactions. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 8359–8367 (2018) [3](#)
26. Goetschalckx, L., Andonian, A., Oliva, A., Isola, P.: Ganalyze: Toward visual definitions of cognitive image properties. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 5744–5753 (2019) [3](#)
27. Goodfellow, I.J., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial networks. arXiv preprint arXiv:1406.2661 (2014) [3](#), [4](#), [6](#)
28. Grabner, H., Gall, J., Van Gool, L.: What makes a chair a chair? In: Proc. IEEE Computer Society Conference on Computer Vision and Pattern Recognition. pp. 1529–1536 (06 2011). <https://doi.org/10.1109/CVPR.2011.5995327> [3](#)
29. Gupta, A., Satkin, S., Efros, A.A., Hebert, M.: From 3d scene geometry to human workspace. In: Computer Vision and Pattern Recognition(CVPR) (2011) [1](#), [3](#)
30. Hassan, M., Ghosh, P., Tesch, J., Tzionas, D., Black, M.J.: Populating 3D scenes by learning human-scene interaction. In: Proceedings IEEE/CVF Conf. on Computer Vision and Pattern Recognition (CVPR) (Jun 2021) [1](#)
31. He, K., Gkioxari, G., Dollár, P., Girshick, R.: Mask r-cnn. In: Proceedings of the IEEE international conference on computer vision. pp. 2961–2969 (2017) [4](#), [14](#)
32. Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., Hochreiter, S.: Gans trained by a two time-scale update rule converge to a local nash equilibrium. In: Proceedings of the 31st International Conference on Neural Information Processing Systems. p. 6629–6640. NIPS’17 (2017) [12](#)
33. Holtzman, A., Buys, J., Du, L., Forbes, M., Choi, Y.: The curious case of neural text degeneration. arXiv preprint arXiv:1904.09751 (2019) [11](#)
34. Härkönen, E., Hertzmann, A., Lehtinen, J., Paris, S.: Ganspace: Discovering interpretable gan controls. In: Proc. NeurIPS (2020) [3](#)
35. Ionescu, C., Papava, D., Olaru, V., Sminchisescu, C.: Human3. 6m: Large scale datasets and predictive methods for 3d human sensing in natural environments. IEEE transactions on pattern analysis and machine intelligence **36**(7), 1325–1339 (2013) [4](#)



36. Isola, P., Zhu, J.Y., Zhou, T., Efros, A.A.: Image-to-image translation with conditional adversarial networks. In: Computer Vision and Pattern Recognition (CVPR), 2017 IEEE Conference on (2017) [3](#), [6](#)
37. Jahanian, A., Chai, L., Isola, P.: On the "steerability" of generative adversarial networks. In: International Conference on Learning Representations (2020) [3](#)
38. Jiang, Y., Koppula, H., Saxena, A.: Hallucinated humans as the hidden context for labeling 3d scenes. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (June 2013) [3](#)
39. Johnson, J., Alahi, A., Fei-Fei, L.: Perceptual losses for real-time style transfer and super-resolution. In: European conference on computer vision. pp. 694–711. Springer (2016) [10](#)
40. Kanazawa, A., Zhang, J.Y., Felsen, P., Malik, J.: Learning 3d human dynamics from video. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5614–5623 (2019) [4](#)
41. Karras, T., Aila, T., Laine, S., Lehtinen, J.: Progressive growing of GANs for improved quality, stability, and variation. In: International Conference on Learning Representations (2018), <https://openreview.net/forum?id=Hk99zCeAb> [3](#), [6](#)
42. Karras, T., Aittala, M., Hellsten, J., Laine, S., Lehtinen, J., Aila, T.: Training generative adversarial networks with limited data. arXiv preprint arXiv:2006.06676 (2020) [6](#), [12](#)
43. Karras, T., Laine, S., Aila, T.: A style-based generator architecture for generative adversarial networks. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4401–4410 (2019) [2](#), [3](#), [6](#), [9](#), [11](#)
44. Karras, T., Laine, S., Aittala, M., Hellsten, J., Lehtinen, J., Aila, T.: Analyzing and improving the image quality of stylegan. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 8110–8119 (2020) [2](#), [3](#), [4](#), [5](#), [6](#), [12](#)
45. Kay, W., Carreira, J., Simonyan, K., Zhang, B., Hillier, C., Vijayanarasimhan, S., Viola, F., Green, T., Back, T., Natsev, P., et al.: The kinetics human action video dataset. arXiv preprint arXiv:1705.06950 (2017) [4](#)
46. Kingma, D.P., Dhariwal, P.: Glow: Generative flow with invertible 1x1 convolutions. In: NeurIPS. pp. 10236–10245 (2018), <http://papers.nips.cc/paper/8224-glow-generative-flow-with-invertible-1x1-convolutions> [11](#)
47. Koppula, H.S., Gupta, R., Saxena, A.: Learning human activities and object affordances from rgb-d videos. The International Journal of Robotics Research **32**(8), 951–970 (2013) [3](#)
48. Lee, J., Chai, J., Reitsma, P.S., Hodgins, J.K., Pollard, N.S.: Interactive control of avatars animated with human motion data. In: Proceedings of the 29th annual conference on Computer graphics and interactive techniques. pp. 491–500 (2002) [3](#)
49. Li, X., Liu, S., Kim, K., Wang, X., Yang, M.H., Kautz, J.: Putting humans in a scene: Learning affordance in 3d indoor environments. In: CVPR (2019) [3](#)
50. Li, Y., Huang, C., Loy, C.C.: Dense intrinsic appearance flow for human pose transfer. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (June 2019) [3](#)
51. Li, Y., Singh, K.K., Ojha, U., Lee, Y.J.: Mixnmatch: Multifactor disentanglement and encoding for conditional image generation. In: CVPR (2020) [3](#)
52. Liu, Z., Luo, P., Wang, X., Tang, X.: Deep learning face attributes in the wild. In: Proceedings of International Conference on Computer Vision (ICCV) (December 2015) [6](#), [9](#)

53. Ma, L., Jia, X., Sun, Q., Schiele, B., Tuytelaars, T., Van Gool, L.: Pose guided person image generation. In: *Advances in Neural Information Processing Systems*. pp. 405–415 (2017) [3](#)
54. Ma, L., Sun, Q., Georgoulis, S., Van Gool, L., Schiele, B., Fritz, M.: Disentangled person image generation. In: *The IEEE International Conference on Computer Vision and Pattern Recognition (CVPR)* (June 2018) [3](#)
55. Van der Maaten, L., Hinton, G.: Visualizing data using t-sne. *Journal of machine learning research* **9**(11) (2008) [11](#)
56. Marchesi, M.: Megapixel size image creation using generative adversarial networks (2017) [11](#)
57. Mescheder, L., Nowozin, S., Geiger, A.: Which training methods for gans do actually converge? In: *International Conference on Machine Learning (ICML)* (2018) [6](#)
58. Mirza, M., Osindero, S.: Conditional generative adversarial nets. *arXiv preprint arXiv:1411.1784* (2014) [3](#), [4](#)
59. Mokady, R., Yarom, M., Tov, O., Lang, O., Cohen-Or, D., Dekel, T., Irani, M., Mosseri, I.: Self-distilled stylegan: Towards generation from internet photos. *arXiv preprint arXiv:2202.12211* (2022) [6](#), [11](#)
60. Monfort, M., Andonian, A., Zhou, B., Ramakrishnan, K., Bargal, S.A., Yan, T., Brown, L., Fan, Q., Gutfreund, D., Vondrick, C., et al.: Moments in time dataset: one million videos for event understanding. *IEEE transactions on pattern analysis and machine intelligence* **42**(2), 502–508 (2019) [4](#)
61. Mottaghi, R., Chen, X., Liu, X., Cho, N.G., Lee, S.W., Fidler, S., Urtasun, R., Yuille, A.: The role of context for object detection and semantic segmentation in the wild. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 891–898 (2014) [3](#)
62. Park, T., Zhu, J.Y., Wang, O., Lu, J., Shechtman, E., Efros, A.A., Zhang, R.: Swapping autoencoder for deep image manipulation. In: *Advances in Neural Information Processing Systems* (2020) [3](#), [6](#)
63. Peebles, W., Peebles, J., Zhu, J.Y., Efros, A.A., Torralba, A.: The hessian penalty: A weak prior for unsupervised disentanglement. In: *Proceedings of European Conference on Computer Vision (ECCV)* (2020) [3](#)
64. Rabinovich, A., Vedaldi, A., Galleguillos, C., Wiewiora, E., Belongie, S.: Objects in context. In: *2007 IEEE 11th International Conference on Computer Vision*. pp. 1–8. IEEE (2007) [3](#)
65. Reed, S., Akata, Z., Yan, X., Logeswaran, L., Schiele, B., Lee, H.: Generative adversarial text to image synthesis. In: *International Conference on Machine Learning*. pp. 1060–1069. PMLR (2016) [6](#)
66. Sauer, A., Schwarz, K., Geiger, A.: Stylegan-xl: Scaling stylegan to large diverse datasets. *arXiv preprint arXiv:2202.00273* (2022) [6](#)
67. Siarohin, A., Sangineto, E., Lathuiliere, S., Sebe, N.: Deformable gans for pose-based human image generation. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 3408–3416 (2018) [3](#), [5](#)
68. Sigurdsson, G.A., Varol, G., Wang, X., Farhadi, A., Laptev, I., Gupta, A.: Hollywood in homes: Crowdsourcing data collection for activity understanding. In: *European Conference on Computer Vision*. pp. 510–526. Springer (2016) [4](#)
69. Szegedy, C., Liu, W., Jia, Y., Sermanet, P., Reed, S., Anguelov, D., Erhan, D., Vanhoucke, V., Rabinovich, A.: Going deeper with convolutions. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 1–9 (2015) [12](#)

70. Torralba, A., Murphy, K.P., Freeman, W.T., Rubin, M.A.: Context-based vision system for place and object recognition. In: Computer Vision, IEEE International Conference on. vol. 2, pp. 273–273. IEEE Computer Society (2003) 3
71. Wang, J., Xu, H., Xu, J., Liu, S., Wang, X.: Synthesizing long-term 3d human motion and interaction in 3d scenes (2020) 3
72. Wang, T.C., Liu, M.Y., Zhu, J.Y., Liu, G., Tao, A., Kautz, J., Catanzaro, B.: Video-to-video synthesis. In: Advances in Neural Information Processing Systems (NeurIPS) (2018) 3
73. Wang, T.C., Liu, M.Y., Zhu, J.Y., Tao, A., Kautz, J., Catanzaro, B.: High-resolution image synthesis and semantic manipulation with conditional gans. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (2018) 6
74. Wang, X., Girdhar, R., Gupta, A.: Binge watching: Scaling affordance learning from sitcoms. In: CVPR (2017) 3
75. Wu, Y., Kirillov, A., Massa, F., Lo, W.Y., Girshick, R.: Detectron2. <https://github.com/facebookresearch/detectron2> (2019) 4, 14
76. Xu, N., Yang, L., Fan, Y., Yang, J., Yue, D., Liang, Y., Price, B., Cohen, S., Huang, T.: Youtube-vos: Sequence-to-sequence video object segmentation. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 585–601 (2018) 4
77. Yao, B., Fei-Fei, L.: Modeling mutual context of object and human pose in human-object interaction activities. 2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition pp. 17–24 (2010) 3
78. Yu, F., Zhang, Y., Song, S., Seff, A., Xiao, J.: Lsun: Construction of a large-scale image dataset using deep learning with humans in the loop. CoRR **abs/1506.03365** (2015), <http://dblp.uni-trier.de/db/journals/corr/corr1506.html#YuZSSX15> 6
79. Zhang, R., Isola, P., Efros, A.A., Shechtman, E., Wang, O.: The unreasonable effectiveness of deep features as a perceptual metric. In: CVPR (2018) 10
80. Zhang, W., Zhu, M., Derpanis, K.G.: From actemes to action: A strongly-supervised representation for detailed action understanding. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 2248–2255 (2013) 4
81. Zhao, S., Liu, Z., Lin, J., Zhu, J.Y., Han, S.: Differentiable augmentation for data-efficient gan training. arXiv preprint arXiv:2006.10738 (2020) 6
82. Zhu, Y., Fathi, A., Fei-Fei, L.: Reasoning about object affordances in a knowledge base representation. In: European conference on computer vision. pp. 408–424. Springer (2014) 3