

# Visual Deprojection: Probabilistic Recovery of Collapsed Dimensions

Guha Balakrishnan  
MIT  
balakg@mit.edu

Adrian V. Dalca  
MIT and MGH  
adalca@mit.edu

Amy Zhao  
MIT  
xamyzhao@mit.edu

John V. Guttag  
MIT  
gutttag@mit.edu

Frédo Durand  
MIT  
fredo@mit.edu

William T. Freeman  
MIT  
freeman@mit.edu

## Abstract

We introduce *visual deprojection*: the task of recovering an image or video that has been collapsed along a dimension. Projections arise in various contexts, such as long-exposure photography, where a dynamic scene is collapsed in time to produce a motion-blurred image, and corner cameras, where reflected light from a scene is collapsed along a spatial dimension because of an edge occluder to yield a 1D video. Deprojection is ill-posed—often there are many plausible solutions for a given input. We first propose a probabilistic model capturing the ambiguity of the task. We then present a variational inference strategy using convolutional neural networks as functional approximators. Sampling from the inference network at test time yields plausible candidates from the distribution of original signals that are consistent with a given input projection. We evaluate the method on several datasets for both spatial and temporal deprojection tasks. We first demonstrate the method can recover human gait videos and face images from spatial projections, and then show that it can recover videos of moving digits from dramatically motion-blurred images obtained via temporal projection.

## 1. Introduction

Captured visual data is often a projection of a higher-dimensional signal “collapsed” along some dimension. For example, long-exposure, motion-blurred photographs are produced by projecting motion trajectories along the time dimension [11, 25]. Recent “corner cameras” leverage the fact that a corner-like edge occluder vertically projects light rays of hidden scenes to produce a 1D video [4]. Medical x-ray machines use spatial projectional radiography, where x-rays are distributed by a generator, and the imaged anatomy affects the signal captured by the detector [26]. Given projected data, is it possible to synthesize the orig-

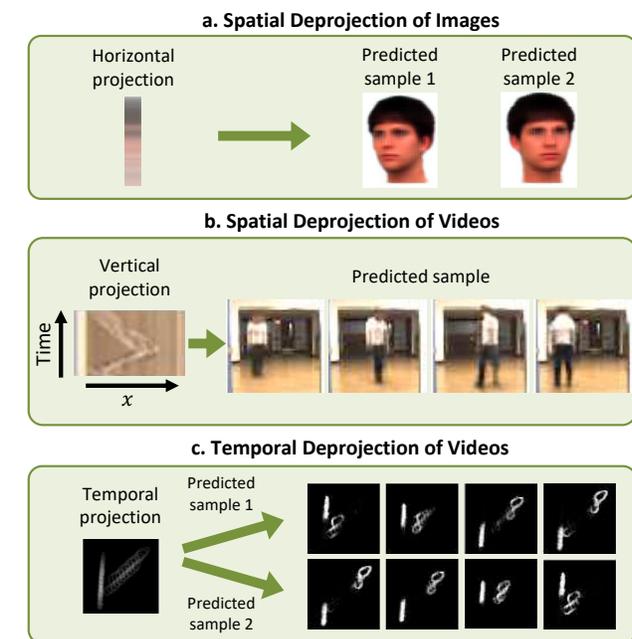


Figure 1: Our method takes a spatial projection of an image or video (a, b) or a temporal projection of a video (c), and outputs a distribution over possible original signals. A projection here is an average of pixel values along a dimension of the original signal. The original signal is only one of multiple possible signals that may have plausibly generated that particular projection.

inal signal? In this work, we present an algorithm that enables this synthesis. We focus on recovering images and video from spatial projections, and recovering a video from a long-exposure image obtained via temporal projection.

The task of inverting projected, high-dimensional signals is ill-posed, making the task infeasible without some priors or constraints on the true signal. This ambiguity includes

object orientations and poses in spatial projections, and the “arrow of time” [43] in temporal projections (Fig. 1). We leverage the fact that the effective dimension of most natural images is often much lower than the pixel representation, because of the shared structure in a given domain. We handle this ambiguity by formulating a probabilistic model for the generation of signals given a projection. The model consists of parametric functions that we implement with convolutional neural networks (CNNs). Using variational inference, we derive an intuitive objective function. Sampling from this deprojection network at test time produces plausible examples of signals that are consistent with an input projection.

There is a rich computer vision literature on recovering high-dimensional data from partial observations. Single-image super-resolution [15], image demosaicing [46], and motion blur removal [14] are all special cases. Here, we focus on projections where a spatial or temporal dimension is entirely removed, resulting in dramatic loss of information. To the best of our knowledge, ours is the first general recovery method in the presence of a collapsed dimension. We build on insights from related problems to develop a first solution for extrapolating appearance and motion cues (in the case of videos) to unseen dimensions. In particular, we leverage recent advances in neural network-based synthesis and stochastic prediction tasks [2, 17, 44].

We evaluate our work both quantitatively and qualitatively. We demonstrate that our method can recover the distribution of human gait videos from 2D spacetime images, and face images from their 1D spatial projections. We also show that our method can model distributions of videos conditioned on motion-blurred images using the Moving MNIST dataset [37].

## 2. Related Work

Projections play a central role in computer vision, starting from the initial stages of image formation, where light from the 3D world is projected onto a 2D plane. We focus on a particular class of projections where higher-dimensional signals of interest are collapsed along one dimension to produce observed data.

### 2.1. Corner Cameras

Corner cameras exploit reflected light from a hidden scene occluded by obstructions with edges to “see around the corner” [4]. Reflected light rays from the scene from the same angular position relative to the corner are vertically integrated to produce a 1D video (one spatial dimension + time). That study used the temporal gradient of the 1D video to coarsely indicate angular positions of the human with respect to the corner, but did not reconstruct the hidden scene. As an initial step towards this difficult reconstruction task, we show that videos and images can be

recovered after collapsing one spatial dimension.

### 2.2. Compressed Sensing

Compressed sensing techniques efficiently reconstruct a signal from limited observations by finding solutions to underdetermined linear systems [8, 12]. This is possible because of the redundancy of natural signals in an appropriate basis. Several methods show that it is possible to accurately reconstruct a signal from a small number (1000s) of bases through convex optimization, even when the bases are chosen randomly [6, 7, 16]. We tackle an extreme variant where one dimension of a signal is completely lost. We also take a learning-based approach to the problem that yields a distribution of potential signals instead of one estimate.

### 2.3. Conditional Image/Video Synthesis and Future Frame Prediction

Neural network-based image and video synthesis has received significant attention. In conditional image synthesis, an image is synthesized conditioned on some other information, such as a class label or another image of the same dimension (image-to-image translation) [5, 17, 29, 38, 42, 47]. In contrast to our work, most of these studies condition on data of the same dimensionality as the output.

Video synthesis algorithms mainly focus on unconditional generation [33, 39, 40] or video-to-video translation [9, 34, 41]. In future video frame prediction, frames are synthesized conditioned on one or more past images. Several of these algorithms treat video generation as a stochastic problem [2, 24, 44], using a variational autoencoder (VAE) style framework [23]. The inputs and outputs in these problems take a similar form to ours, but the information in the input is different. We draw insights from the stochastic formulation in these studies for our task.

### 2.4. Inverting a Motion-blurred Image to Video

One application we explore is the formation of videos from dramatically motion-blurred images, created by temporally aggregating photons from a scene over an extended period of time. Two recent studies present the deterministic recovery of a video sequence from a single motion-blurred image [18, 30]. We propose a general deprojection framework for dimensions including, but not limited to time. In addition, our framework is probabilistic, capturing the distribution of signal variability instead of a single deterministic output (see Fig. 1).

## 3. Methods

We assume a dataset of pairs  $\{\mathbf{x}, \mathbf{y}\}$  of original signals  $\mathbf{y} \in \mathbb{R}^{d_1 \times \dots \times d_D}$  and projections  $\mathbf{x} \in \mathbb{R}^{d_1 \times \dots \times d_{p-1} \times d_{p+1} \times \dots \times d_D}$ , where  $D$  is the number of dimensions of  $\mathbf{y}$  and  $p$  is the projected dimension.

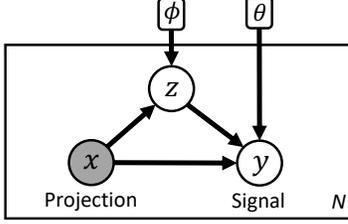


Figure 2: Schematic of the probabilistic model at test time. The shaded variable  $\mathbf{x}$  is the observed input projection,  $\mathbf{y}$  is the higher-dimensional signal,  $\mathbf{z}$  is a multinomial latent variable,  $\phi$  and  $\theta$  are global network parameters, and  $N$  is the number of test examples in the dataset.

We assume a projection function  $f_\omega : \mathbb{R}^D \rightarrow \mathbb{R}^{D-1}$  with parameters  $\omega$ . In our experiments, we focus on a case often observed in practice, where  $f_\omega$  is a linear operation in  $\omega$  along  $p$ , such as averaging:  $\mathbf{x} = f_\omega(\mathbf{y}) = \sum_{k=1}^{d_p} \omega_k \mathbf{y}^{p=k}$ , where  $\mathbf{y}^{p=k}$  is the  $k^{\text{th}}$  slice of  $\mathbf{y}$  along dimension  $p$ . For example, a grayscale video  $\mathbf{y} \in \mathbb{R}^{H \times W \times T}$  might get projected to an image  $\mathbf{x} \in \mathbb{R}^{H \times W}$  by averaging pixels across time. Deprojection is a highly underconstrained problem. Even if the values of  $\omega$  are known, there are  $d_p$  as many variables (size of  $\mathbf{y}$ ) as constraints (size of  $\mathbf{x}$ ).

We aim to capture the distribution  $p(\mathbf{y}|\mathbf{x})$  for a particular scenario with data. We first present a probabilistic model for the deprojection task which builds on the conditional VAE (CVAE) [36] (Fig. 2). We let  $\mathbf{z} \sim p_\phi(\mathbf{z}|\mathbf{x})$  be a multivariate normal latent variable which captures variability of  $\mathbf{y}$  unexplainable from  $\mathbf{x}$  alone. Intuitively,  $\mathbf{z}$  encodes information orthogonal to the unprojected dimensions. For example, it could capture the temporal variation of the various scenes that may have led to a long-exposure image.

We define  $p_\theta(\mathbf{y}|\mathbf{x}, \mathbf{z})$  as a Gaussian distribution:

$$p_\theta(\mathbf{y}|\mathbf{x}, \mathbf{z}) = \mathcal{N}(\mathbf{y}; g_\theta(\mathbf{x}, \mathbf{z}), \mathbf{I}\sigma_y^2) \quad (1)$$

where  $\sigma_y^2$  is a per-pixel noise variance and  $g_\theta(\mathbf{x}, \mathbf{z})$  is a *deprojection* function, parameterized by  $\theta$  and responsible for producing a noiseless estimate of  $\mathbf{y}$  given  $\mathbf{x}$  and  $\mathbf{z}$ .

### 3.1. Variational Inference and Loss Function

Our goal is to estimate  $p_{\phi, \theta}(\mathbf{y}|\mathbf{x})$ :

$$p_{\phi, \theta}(\mathbf{y}|\mathbf{x}) = \int_{\mathbf{z}} p_\theta(\mathbf{y}|\mathbf{x}, \mathbf{z}) p_\phi(\mathbf{z}|\mathbf{x}) d\mathbf{z} \quad (2)$$

Evaluating this integral directly is intractable because of its reliance on potentially complex parametric functions and the intractability of estimating the posterior  $p(\mathbf{z}|\mathbf{y})$ . We instead use variational inference to obtain a lower bound of

the likelihood, and use stochastic gradient descent to optimize it [20, 23]. We introduce an approximative posterior distribution  $q_\psi(\mathbf{z}|\mathbf{y}) = \mathcal{N}(\mathbf{z}; \mu_\psi(\mathbf{y}), \sigma_\psi(\mathbf{y}))$ :

$$\log p_{\phi, \theta}(\mathbf{y}|\mathbf{x}) = \log E_{\mathbf{z} \sim q_\psi} \left[ \frac{p_\phi(\mathbf{z}|\mathbf{x})}{q_\psi(\mathbf{z}|\mathbf{y})} p_\theta(\mathbf{y}|\mathbf{x}, \mathbf{z}) \right]. \quad (3)$$

Using Jensen’s inequality, we achieve the following evidence lower bound (ELBO) for  $\log p_{\phi, \theta}(\mathbf{y}|\mathbf{x})$ :

$$\log p_{\phi, \theta}(\mathbf{y}|\mathbf{x}) \geq E_{\mathbf{z} \sim q_\psi} [\log p_\theta(\mathbf{y}|\mathbf{x}, \mathbf{z})] - D_{KL}[q_\psi(\mathbf{z}|\mathbf{y}) || p_\phi(\mathbf{z}|\mathbf{x})], \quad (4)$$

where  $D_{KL}[\cdot || \cdot]$  is the Kullback-Leibler divergence encouraging the variational distribution to approximate the conditional prior, resulting in a regularized embedding. We estimate the expectation term by drawing one  $\hat{\mathbf{z}}$  from  $q_\psi(\mathbf{z}|\mathbf{y})$  within the network using the reparametrization trick [23] and evaluating the expression:

$$\log p_\theta(\mathbf{y}|\mathbf{x}, \hat{\mathbf{z}}) = \frac{\|g_\theta(\mathbf{x}, \hat{\mathbf{z}}) - \mathbf{y}\|_2^2}{2\sigma_y^2} + \text{const.} \quad (5)$$

This leads to the training loss function to be minimized:

$$\mathcal{L}_{\phi, \psi, \theta}(\mathbf{x}, \mathbf{y}, \hat{\mathbf{z}}) = \beta D_{KL}[q_\psi(\mathbf{z}|\mathbf{y}) || p_\phi(\mathbf{z}|\mathbf{x})] + \|g_\theta(\mathbf{x}, \hat{\mathbf{z}}) - \mathbf{y}\|_2^2 \quad (6)$$

where  $\beta$  is a tradeoff parameter capturing the relative importance of the regularization term. The per-pixel reconstruction term in Eq. (6) can result in blurry outputs. For datasets with subtle details such as face images, we also add a perceptual error, computed over a learned feature space [13, 19, 45]. We use a distance function  $D_\gamma(\cdot, \cdot)$  [45], computed over high-dimensional features learned by the VGG16 network [35] with parameters  $\gamma$ , trained to perform classification on ImageNet.

## 3.2. Network Architectures

We implement  $g_\theta(\cdot, \cdot)$  and the Gaussian parameters of  $q_\psi(\cdot|\cdot)$  and  $p_\phi(\cdot|\cdot)$  with neural networks. Fig. 3 depicts the architecture for the 2D-to-3D temporal deprojection task. Our 2D-to-3D spatial deprojection architecture is nearly identical, differing only in the dimensions of  $\mathbf{x}$  and the reshaping operator’s dimension ordering. We handle 1D-to-2D deprojections by using the lower-dimensional versions of the convolution and reshaping operators. The number of convolutional layers, and number of parameters vary by dataset based on their complexities.

### 3.2.1 Posterior and Prior Encoders

The encoder for the distribution parameters of the posterior  $q_\psi(\cdot|\cdot)$  is implemented using a series of strided 3D convolutional operators and Leaky ReLU activations until a volume

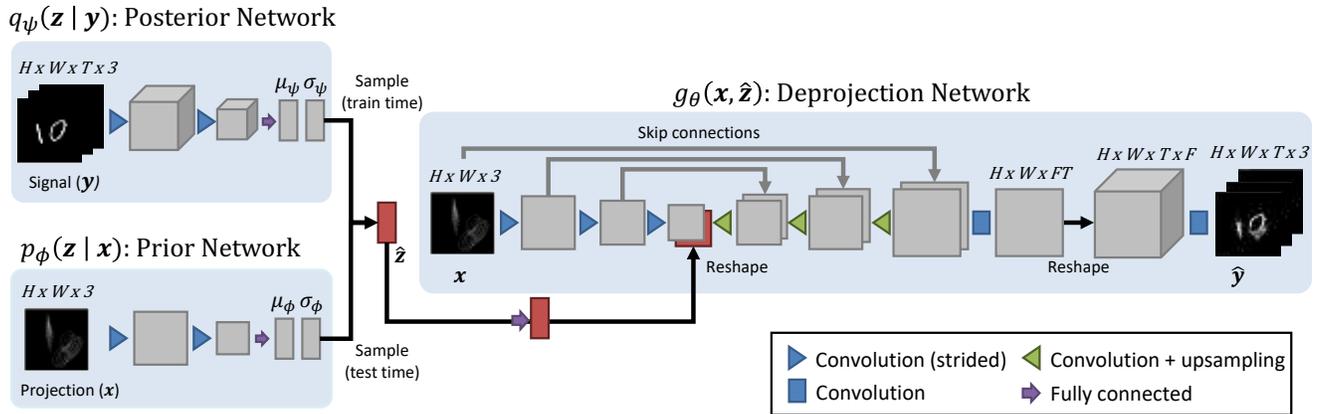


Figure 3: Overview of our network architecture, drawn here for the 2D-to-3D temporal deprojection task. The network consists of three parameterized functions:  $q_\psi(\cdot|\cdot)$  captures the variational posterior distribution,  $p_\phi(\cdot|\cdot)$  captures the prior distribution and  $g_\theta(\cdot, \cdot)$  performs deprojection.  $\mathbf{z}$  is sampled from  $q_\psi(\cdot)$  during training, and from  $p_\phi(\cdot)$  during testing.

of resolution less than  $8 \times 8 \times 3$  is reached. We flatten this volume and use two fully connected layers to obtain  $\mu_\psi$  and  $\sigma_\psi$ , the distribution parameters. The encoder for the conditional prior  $p_\phi(\cdot|\cdot)$  is implemented in a similar way, with 2D strided convolutions. One  $\hat{\mathbf{z}}$  is drawn from  $q_\psi(\cdot|\cdot)$  and fed to the deprojection function. At test time,  $\hat{\mathbf{z}}$  is drawn from  $p_\phi(\cdot|\cdot)$  to visualize results.

### 3.2.2 Deprojection Function

The function  $g_\theta(\mathbf{x}, \hat{\mathbf{z}})$  deprojects  $\mathbf{x}$  into an estimate  $\hat{\mathbf{y}}$ . We first use a UNet-style architecture [32] to compute per-pixel features of  $\mathbf{x}$ . The UNet consists of two stages. In the first stage, we apply a series of strided 2D convolutional operators to extract multiscale features. We apply a fully connected layer to  $\hat{\mathbf{z}}$ , reshape these activations into an image, and concatenate this image to the coarsest features. The second stage applies a series of 2D convolutions and upsampling operations to synthesize an image of the same dimensions as  $\mathbf{x}$  and many more data channels. Activations from the first stage are concatenated to the second stage activations with skip connections to propagate learned features.

We expand the resulting image along the collapsed dimension to produce a 3D volume. To do this, we apply a 2D convolution to produce  $TF$  data channels, where  $T$  is the size of the collapsed dimension (time in this case), and  $F$  is some number of features. Finally, we reshape this image into a 3D volume, and apply a few 3D convolutions to refine and produce a signal estimate  $\hat{\mathbf{y}}$ .

## 4. Experiments and Results

We first evaluate our method on 1D-to-2D spatial deprojections of human faces using FacePlace [31]. We then

show results for 2D-to-3D spatial deprojections using an in-house dataset of human gait videos collected by the authors. Finally, we demonstrate 2D-to-3D temporal deprojections using the Moving MNIST [37] dataset. We focus on projections where pixels are averaged along a dimension for all experiments. For all experiments we split the data into train/test/validation non-overlapping groups.

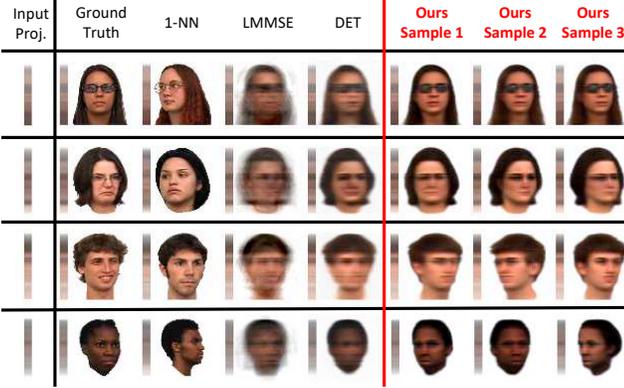
### 4.1. Implementation

We implement our models in Keras [10] with a TensorFlow [1] backend. We use the ADAM optimizer [22] with a learning rate of  $1e^{-4}$ . We trained separate models for each experiment. We select the regularization hyperparameter  $\beta$  separately for each dataset such that the KL term is between [5, 15] on our validation data, to obtain adequate data reconstruction while avoiding mode collapse. We set the dimension of  $\mathbf{z}$  to 10 for all experiments.

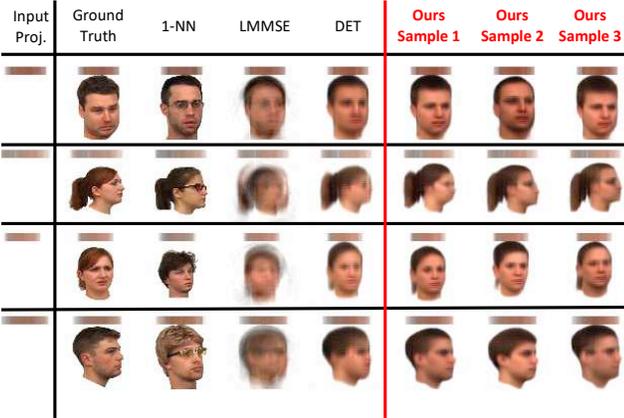
### 4.2. Spatial Deprojections with FacePlace

FacePlace consists of over 5,000 images of 236 different people. There are many sources of variability, including different ethnicities, multiple views, facial expressions, and props. We randomly held out all images for 30 individuals to form a test set. We scaled images to  $128 \times 128$  pixels and performed data augmentation with translation, scaling and saturation variations. We compare our method against the following baselines:

1. Nearest neighbor selector ( $k$ -NN): Selects the  $k$  images from the training dataset with projections closest to the test projection using mean squared error distance.
2. A deterministic model (DET) identical to the deprojection network  $g_\theta(\mathbf{x}, \mathbf{z})$  of our method, without the



(a) Horizontal projection results.



(b) Vertical projection results.

Figure 4: Sample image reconstructions on FacePlace. The input projections along with ground truth images are shown on the left. Our method’s samples are randomly chosen. Our method is able to synthesize a variety of appearances with projections closely matching the input.

incorporation of a latent variable  $\mathbf{z}$ .

3. A linear minimum mean squared error (LMMSE) estimator which assumes that  $\mathbf{x}$  and  $\mathbf{y}$  are drawn from distributions  $X, Y$  such that  $\bar{\mathbf{y}} = E_Y[\mathbf{y}]$  is linear in  $\mathbf{x}$ :  $\bar{\mathbf{y}} = A\mathbf{x} + b$  for some parameters  $A$  and  $b$ . Minimizing the expected MSE of  $\mathbf{y}$  yields a closed form expression for  $p(\mathbf{y}|\mathbf{x})$ :

$$p(\mathbf{y}|\mathbf{x}) = \mathcal{N}(\mathbf{y}; \Sigma_{YX}\Sigma_X^{-1}(\mathbf{x} - \bar{\mathbf{x}}) + \bar{\mathbf{y}}, \Sigma_Y - \Sigma_{YX}\Sigma_X^{-1}\Sigma_{YX}^T), \quad (7)$$

where  $\Sigma_X$  and  $\Sigma_Y$  are the covariance matrices of  $X$  and  $Y$  and  $\Sigma_{XY}$  is their cross-covariance matrix.

For both our method and DET, we used the perceptual loss metric. Fig. 4 presents visual results, with a few randomly chosen samples from our method. 1-NN varies in

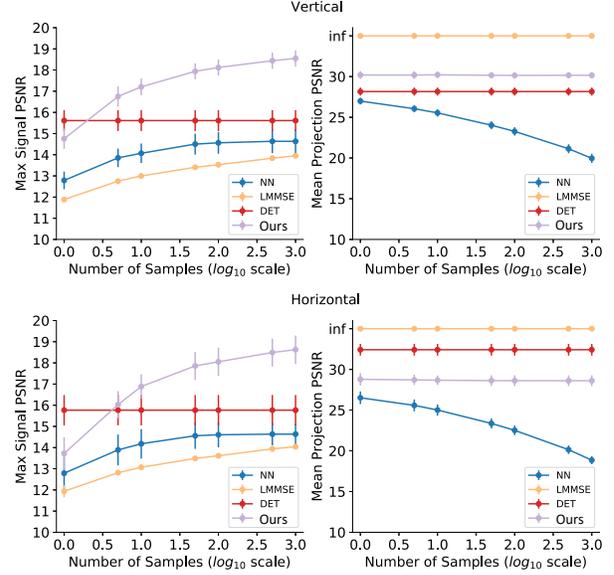


Figure 5: FacePlace PSNR for all methods (vertical projection on top, horizontal on bottom, max signal PSNR (deprojection estimate) on left, mean projection PSNR on right) with varying sample size for 100 test projections. Our method yields higher maximum signal PSNR than all baselines. DET has a higher expected signal PSNR for one sample because it tends to return a blurry average over many signals. LMMSE has infinite projection PSNR because it captures the exact linear signal-projection relationship by construction.

performance depending on the test example, and can sometimes produce faces from the wrong person. LMMSE produces very blurry outputs, indicating the highly nonlinear nature of this task. DET produces less blurry outputs, but still often merges different plausible faces together. Our method captures uncertainty of head orientations as well as appearance variations, such as hair color and facial structure. Ambiguity in head orientation is more apparent with the horizontal projections, since pose changes affect that dimension the most. The outputs of our methods are also sharper than LMMSE and DET, and are more consistent with ground truth than 1-NN.

We also quantitatively evaluate the models. We use PSNR (peak-signal-to-noise-ratio, higher is better) to measure reconstruction quality between images. For each test projection, we sample  $k$  deprojection estimates from each model (DET always returns the same estimate) and record the highest PSNR between any estimate and the ground truth image. For each deprojection estimate, we reproject and record the average PSNR of the output projections with respect to the the test (initial) projection.

Fig. 5 illustrates the results with varying samples  $k$  for

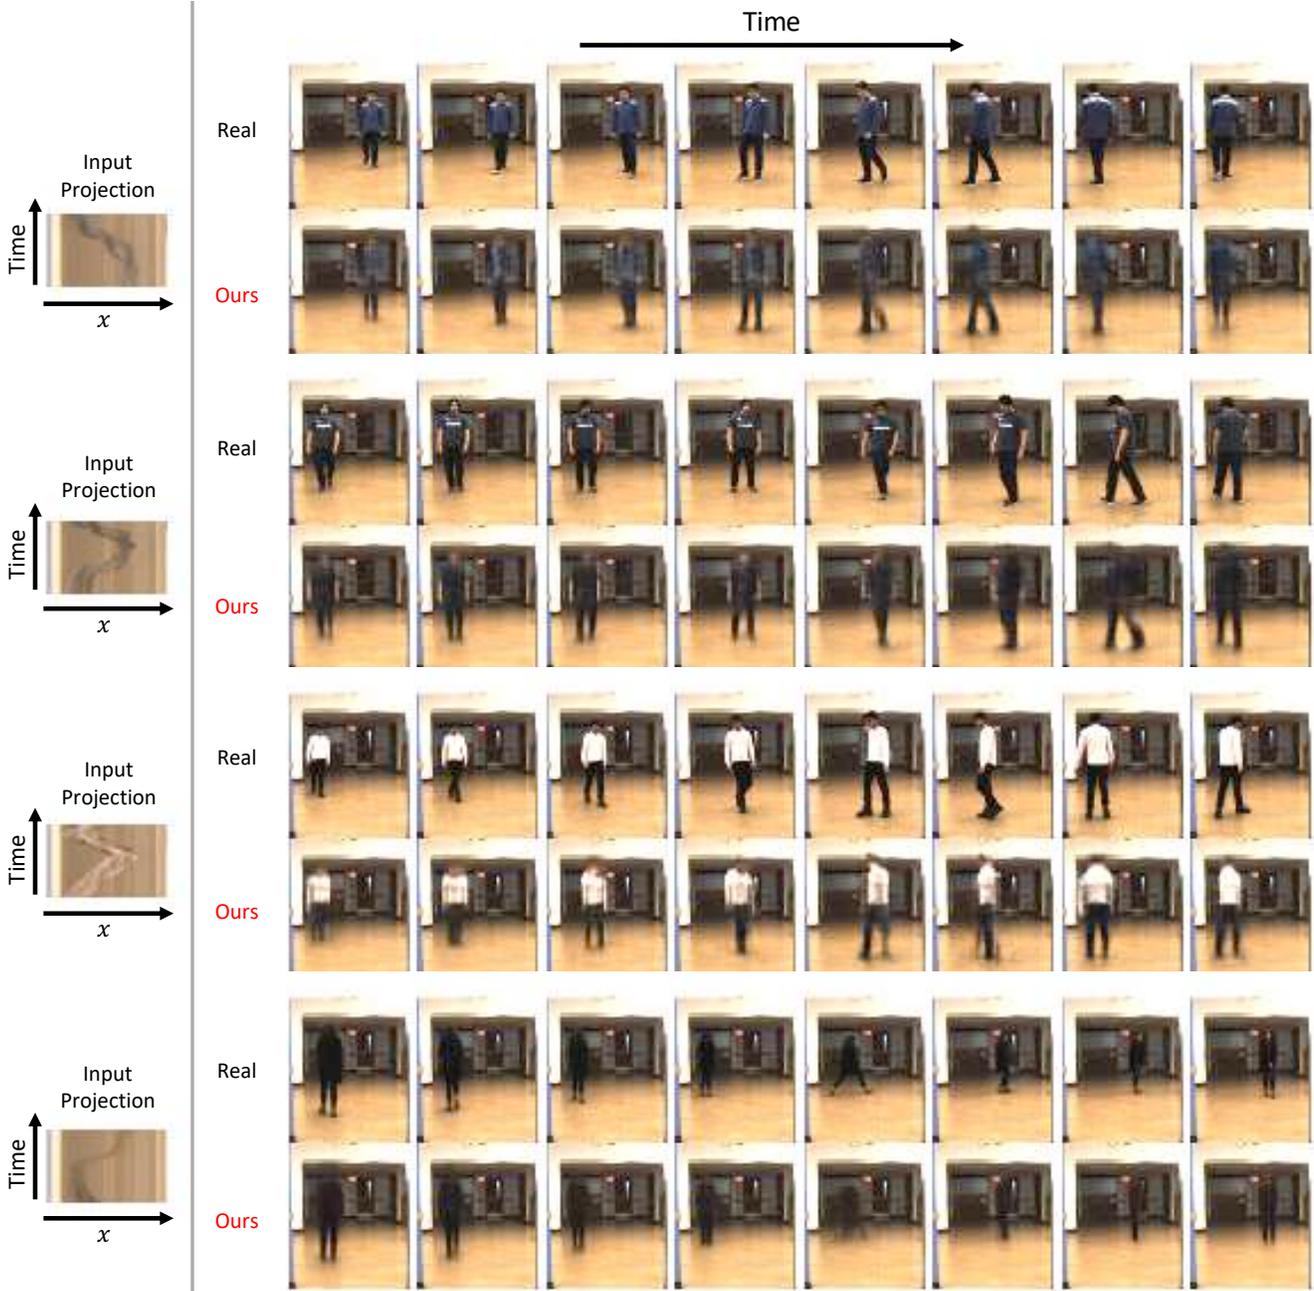


Figure 6: Sample outputs for four examples from the in-house walking dataset. The left column shows the input vertical projection. For each example, the top row displays the ground truth sequence and the bottom row displays our method’s mean output using  $\mathbf{z} = \mu_\phi$ .

100 test projections. As the number of samples  $k$  increases, our method’s signal (deprojection) PSNR improves, highlighting the advantage of our probabilistic approach. Best estimates from  $k$ -NN approach the best estimates of our method in signal reconstruction with increasing  $k$ , but many poor estimates are also retrieved by  $k$ -NN as evidenced by its decreasing projection PSNR curve. LMMSE has perfect

projection PSNR because it captures the exact linear relationship between the signal and projection by construction. DET has higher signal PSNR when drawing one sample, because it averages over plausible images, while our method does not. Our proposed method surpasses DET after 1 sample.

### 4.3. Spatial Deprojections with Walking Videos

We qualitatively evaluate our method on reconstructing human gait videos from vertical spatial projections. This scenario is of practical relevance for corner cameras, described in Sec. 2.1. We collected 35 videos of 30 subjects walking in a specified area for one minute each. Subjects had varying attire, heights (5'2"- 6'5"), ages (18-60), and sexes (18m/12f). Subjects were not instructed to walk in any particular way, and many walked in odd patterns. The background is identical for all videos. We downsampled the videos to 5 frames per second and each frame to  $256 \times 224$  pixels, and apply data augmentation of horizontal translations to each video. We held out 6 subjects to produce a test set. We predict sequences of 24 frames (roughly 5 seconds in real time).

Fig. 6 presents several reconstruction examples, obtained by setting  $\hat{z} = \mu_\phi(\mathbf{x})$ , the mean of the prior distribution. Our method recovers many details from the vertical projections alone. The background is easily synthesized because it is consistent among all videos in the dataset. Remarkably, many appearance and pose details of the subjects are also recovered. Subtle fluctuations in pixel intensity and the shape of the projected foreground trace contain clues about the foreground signal along the collapsed dimension. For example, the method seems to learn that a trace that gets darker and wider with time likely corresponds to a person walking closer to the camera.

The third subject is an illustrative result for which our method separates the white shirt from black pants despite their aspects not being obvious in the projection. Projected details, along with a learned pattern that shirts are often lighter colors than pants, likely enable this recovery. Finally, the method may struggle with patterns rarely seen in the training data, such as the large step by the fourth subject in the fifth frame.

In addition to these experiments, we trained a separate model on the DGAIT dataset [3] consisting of more subjects (53), but with simpler walking patterns. We obtain results with similar quality, as illustrated in Fig. 7.

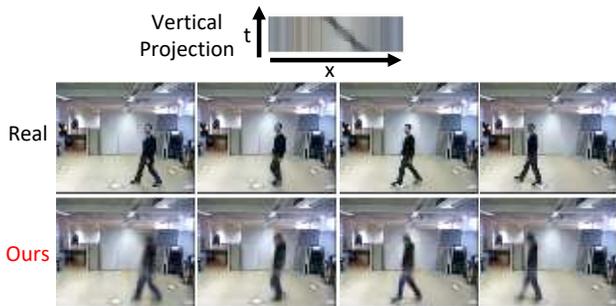


Figure 7: Sample output from the DGAIT walking dataset.

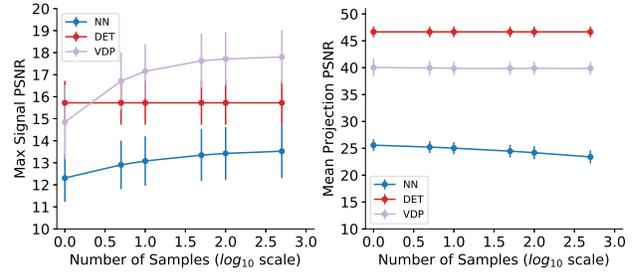


Figure 8: Moving MNIST PSNR plots for 100 projection test examples, similar to Fig. 5.

### 4.4. Temporal Deprojections with Moving MNIST

The Moving MNIST dataset consists of 10,000 video sequences of two moving handwritten digits. The digits can occlude one another, and bounce off the edges of the frames. Given a dataset of  $64 \times 64 \times 10$ -sized video subclips, we generate each projection  $\mathbf{x}$  by averaging the frames in time, similar to other studies that generate motion-blurred images at a large scale [18, 21, 27, 28]. Despite the simple appearance and dynamics of this dataset, synthesizing digit appearances and capturing the plausible directions of each trajectory is challenging.

Sample outputs of our method for three test examples are visualized in Fig. 9. To illustrate the temporal aspects learned by our method, we sample 10 sequences from our method for each projection, and present the sequences with the lowest mean squared error with respect to the ground truth clip run forwards and backwards. Our method is able to infer the shape of the characters from a dramatically motion-blurred input image, difficult to interpret even by human standards. Furthermore, our method captures the multimodal dynamics of the dataset, which we illustrate by presenting the two motion sequences: the first sequence matches the temporal direction of the ground truth, and the second matches the reverse temporal progression.

We quantify our accuracy using PSNR curves, similar to the first experiment, displayed in Fig. 8. Because of the prohibitive computational costs of generating the full joint covariance matrix, we do not evaluate LMMSE in this experiment. DET produces blurry sequences, by merging different plausible temporal orderings. Similar to the first experiment, this results in DET outputs having the best expected signal (deprojection) PSNR only for  $k = 1$ . Our method clearly outperforms DET in signal PSNR for  $k > 1$ . DET performs better in projection PSNR, since in this experiment an average of all plausible sequences yields a very accurate projection.  $k$ -NN performs relatively worse in this experiment compared to the FacePlace experiments, because of the difficulty in finding nearest neighbors in higher-dimensions.

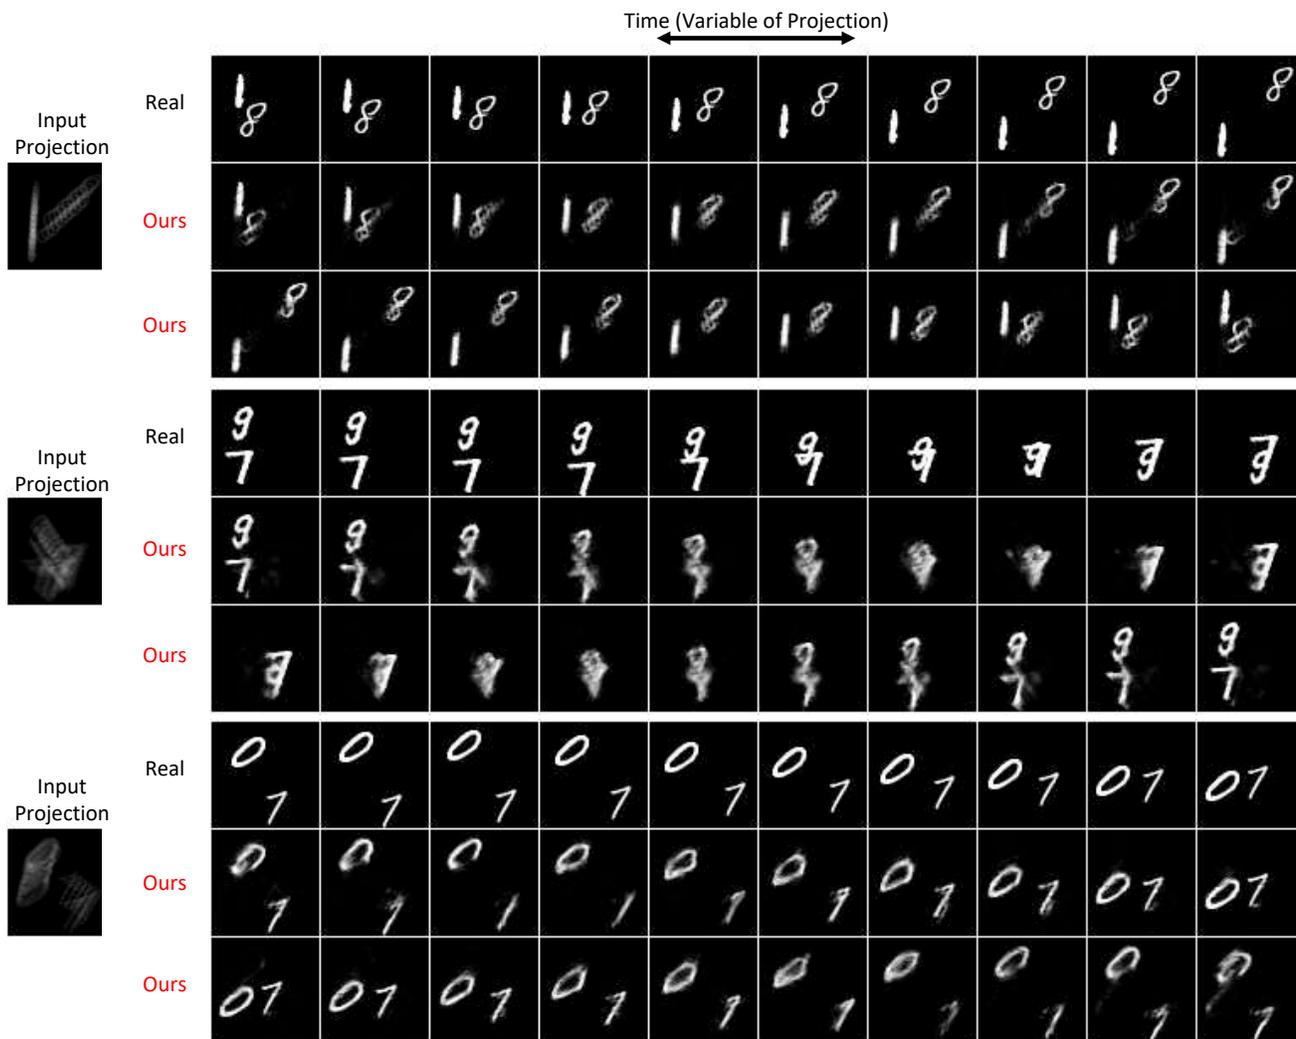


Figure 9: Sample outputs from the Moving MNIST dataset. The left column shows the input projection. For each example, the top row displays the ground truth sequence. We show two sample sequences produced by our method per input projection: the first matches the temporal direction of the ground truth, and the second synthesizes the reverse temporal progression.

## 5. Conclusion

In this work, we introduced the novel problem of *visual deprojection*: synthesizing an image or video that has been collapsed along a dimension into a lower-dimensional observation. We presented a first general method that handles both images and video, and projections along any dimension of these data. We addressed the uncertainty of the task by first introducing a probabilistic model that captures the distribution of original signals conditioned on a projection. We implemented the parameterized functions of this model with CNNs to learn shared image structures in each domain and enable accurate signal synthesis.

Though information from a collapsed dimension is often seemingly unrecoverable from a projection to the naked

eye, our results demonstrate that much of the “lost” information is recoverable. We demonstrated this by reconstructing subtle details of faces in images and accurate motion in videos from spatial projections alone. Finally, we illustrate that videos can be reconstructed from dramatically motion blurred images, even with multimodal trajectories, using the Moving MNIST dataset. This work illustrates promising results in a new, ambitious imaging task and opens exciting possibilities in future applications of revealing the invisible.

### Acknowledgments

This work was funded by DARPA REVEAL Program under Contract No. HR0011-16-C-0030, NIH 1R21AG050122 and Wistron Corp.

## References

- [1] Martín Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, et al. Tensorflow: A system for large-scale machine learning. In *12th {USENIX} Symposium on Operating Systems Design and Implementation ({OSDI} 16)*, pages 265–283, 2016. 4
- [2] Mohammad Babaeizadeh, Chelsea Finn, Dumitru Erhan, Roy H Campbell, and Sergey Levine. Stochastic variational video prediction. *arXiv preprint arXiv:1710.11252*, 2017. 2
- [3] Ricard Borràs, Àgata Lapedriza, and Laura Igual. Depth information in human gait analysis: An experimental study on gender recognition. In *International Conference Image Analysis and Recognition*, pages 98–105. Springer, 2012. 7
- [4] Katherine L Bouman, Vickie Ye, Adam B Yedidia, Frédo Durand, Gregory W Wornell, Antonio Torralba, and William T Freeman. Turning corners into cameras: Principles and methods. In *International Conference on Computer Vision*, volume 1, page 8, 2017. 1, 2
- [5] Konstantinos Bousmalis, Nathan Silberman, David Dohan, Dumitru Erhan, and Dilip Krishnan. Unsupervised pixel-level domain adaptation with generative adversarial networks. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, volume 1, page 7, 2017. 2
- [6] Emmanuel J Candes and Justin K Romberg. Signal recovery from random projections. In *Computational Imaging III*, volume 5674, pages 76–87. International Society for Optics and Photonics, 2005. 2
- [7] Emmanuel J Candes and Terence Tao. Near-optimal signal recovery from random projections: Universal encoding strategies? *IEEE transactions on information theory*, 52(12):5406–5425, 2006. 2
- [8] Emmanuel J Candès and Michael B Wakin. An introduction to compressive sampling. *IEEE signal processing magazine*, 25(2):21–30, 2008. 2
- [9] Dongdong Chen, Jing Liao, Lu Yuan, Nenghai Yu, and Gang Hua. Coherent online video style transfer. In *Proc. Intl. Conf. Computer Vision (ICCV)*, 2017. 2
- [10] François Chollet et al. Keras. <https://keras.io>, 2015. 4
- [11] Shengyang Dai and Ying Wu. Motion from blur. 2008. 1
- [12] David L Donoho. Compressed sensing. *IEEE Transactions on information theory*, 52(4):1289–1306, 2006. 2
- [13] Alexey Dosovitskiy and Thomas Brox. Generating images with perceptual similarity metrics based on deep networks. In *Advances in Neural Information Processing Systems*, pages 658–666, 2016. 3
- [14] Rob Fergus, Barun Singh, Aaron Hertzmann, Sam T Roweis, and William T Freeman. Removing camera shake from a single photograph. In *ACM transactions on graphics (TOG)*, volume 25, pages 787–794. ACM, 2006. 2
- [15] William T Freeman, Thouis R Jones, and Egon C Pasztor. Example-based super-resolution. *IEEE Computer graphics and Applications*, 22(2):56–65, 2002. 2
- [16] Jarvis Haupt and Robert Nowak. Signal reconstruction from noisy random projections. *IEEE Transactions on Information Theory*, 52(9):4036–4048, 2006. 2
- [17] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. *arXiv preprint*, 2017. 2
- [18] Meiguang Jin, Givi Meishvili, and Paolo Favaro. Learning to extract a video sequence from a single motion-blurred image. *arXiv preprint arXiv:1804.04065*, 2018. 2, 7
- [19] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *European Conference on Computer Vision*, pages 694–711. Springer, 2016. 3
- [20] Michael I Jordan, Zoubin Ghahramani, Tommi S Jaakkola, and Lawrence K Saul. An introduction to variational methods for graphical models. *Machine learning*, 37(2):183–233, 1999. 3
- [21] Tae Hyun Kim, Kyoung Mu Lee, Bernhard Schölkopf, and Michael Hirsch. Online video deblurring via dynamic temporal blending network. In *Computer Vision (ICCV), 2017 IEEE International Conference on*, pages 4058–4067. IEEE, 2017. 7
- [22] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 4
- [23] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013. 2, 3
- [24] Alex X Lee, Richard Zhang, Frederik Ebert, Pieter Abbeel, Chelsea Finn, and Sergey Levine. Stochastic adversarial video prediction. *arXiv preprint arXiv:1804.01523*, 2018. 2
- [25] Anat Levin, Peter Sand, Taeg Sang Cho, Fredo Durand, and William T. Freeman. Motion-invariant photography. *ACM Transactions on Graphics (SIGGRAPH 2008)*. 1
- [26] Clint M Logan, Daniel J Schneberk, Peter J Shull, and Harry E Martz. *X-ray Imaging: fundamentals, industrial techniques and applications*. CRC Press, 2016. 1
- [27] Seungjun Nah, Tae Hyun Kim, and Kyoung Mu Lee. Deep multi-scale convolutional neural network for dynamic scene deblurring. In *CVPR*, volume 1, page 3, 2017. 7
- [28] Mehdi Noroozi, Paramanand Chandramouli, and Paolo Favaro. Motion deblurring in the wild. In *German Conference on Pattern Recognition*, pages 65–77. Springer, 2017. 7
- [29] Augustus Odena, Christopher Olah, and Jonathon Shlens. Conditional image synthesis with auxiliary classifier gans. *arXiv preprint arXiv:1610.09585*, 2016. 2
- [30] Kuldeep Purohit, Anshul Shah, and AN Rajagopalan. Bringing alive blurred moments. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6830–6839, 2019. 2
- [31] Giulia Righi, Jessie J Peissig, and Michael J Tarr. Recognizing disguised faces. *Visual Cognition*, 20(2):143–169, 2012. 4
- [32] Olaf Ronneberger, Philipp Fischer, and Thomas 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. 4

- [33] Masaki Saito, Eiichi Matsumoto, and Shunta Saito. Temporal generative adversarial nets with singular value clipping. In *IEEE International Conference on Computer Vision (ICCV)*, volume 2, page 5, 2017. 2
- [34] Eli Shechtman, Yaron Caspi, and Michal Irani. Space-time super-resolution. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 27(4):531–545, 2005. 2
- [35] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. 3
- [36] Kihyuk Sohn, Honglak Lee, and Xinchen Yan. Learning structured output representation using deep conditional generative models. In *Advances in Neural Information Processing Systems*, pages 3483–3491, 2015. 3
- [37] Nitish Srivastava, Elman Mansimov, and Ruslan Salakhudinov. Unsupervised learning of video representations using lstms. In *International conference on machine learning*, pages 843–852, 2015. 2, 4
- [38] Yaniv Taigman, Adam Polyak, and Lior Wolf. Unsupervised cross-domain image generation. *arXiv preprint arXiv:1611.02200*, 2016. 2
- [39] Sergey Tulyakov, Ming-Yu Liu, Xiaodong Yang, and Jan Kautz. Mocogan: Decomposing motion and content for video generation. *arXiv preprint arXiv:1707.04993*, 2017. 2
- [40] Carl Vondrick, Hamed Pirsiavash, and Antonio Torralba. Generating videos with scene dynamics. In *Advances In Neural Information Processing Systems*, pages 613–621, 2016. 2
- [41] Ting-Chun Wang, Ming-Yu Liu, Jun-Yan Zhu, Guilin Liu, Andrew Tao, Jan Kautz, and Bryan Catanzaro. Video-to-video synthesis. In *Advances in Neural Information Processing Systems (NIPS)*, 2018. 2
- [42] Ting-Chun Wang, Ming-Yu Liu, Jun-Yan Zhu, Andrew Tao, Jan Kautz, and Bryan Catanzaro. High-resolution image synthesis and semantic manipulation with conditional gans. *arXiv preprint arXiv:1711.11585*, 2017. 2
- [43] Donglai Wei, Joseph J Lim, Andrew Zisserman, and William T Freeman. Learning and using the arrow of time. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8052–8060, 2018. 2
- [44] Tianfan Xue, Jiajun Wu, Katherine Bouman, and William Freeman. Visual dynamics: Stochastic future generation via layered cross convolutional networks. *IEEE transactions on pattern analysis and machine intelligence*, 2018. 2
- [45] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. *arXiv preprint*, 2018. 3
- [46] Ruiwen Zhen and Robert L Stevenson. Image demosaicing. In *Color image and video enhancement*, pages 13–54. Springer, 2015. 2
- [47] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. *arXiv preprint*, 2017. 2