Comparing Correspondences: Video Prediction with Correspondence-wise Losses

Daniel Geng   Max Hamilton   Andrew Owens

University of Michigan

Abstract

Image prediction methods often struggle on tasks that require changing the positions of objects, such as video prediction, producing blurry images that average over the many positions that objects might occupy. In this paper, we propose a simple change to existing image similarity metrics that makes them more robust to positional errors: we match the images using optical flow, then measure the visual similarity of corresponding pixels. This change leads to crisper and more perceptually accurate predictions, and does not require modifications to the image prediction network. We apply our method to a variety of video prediction tasks, where it obtains strong performance with simple network architectures, and to the closely related task of video interpolation. Code and results are available at our webpage: https://dangeng.github.io/CorrWiseLosses

1. Introduction

Recent years have seen major advances in image prediction [6, 11, 38, 46, 71], yet these methods often struggle to successfully alter *image structure*. Consequently, tasks that involve modifying the positions or shapes of objects, such as video prediction and the closely related problem of video interpolation, remain challenging open problems.

Often, there is fundamental uncertainty over where exactly an object should be. When this happens, models tend to produce blurry results. This undesirable behavior is often encouraged by the loss function. Under simple *pixel-wise* loss functions, such as the L_1 distance, each incorrectly positioned pixel is compared to a pixel that belongs to a different object, thereby incurring a large penalty. Models trained using these losses therefore "hedge" by averaging over all of the possible positions an object might occupy, resulting in images with significantly lower loss.

We take inspiration from classic image matching methods, such as Hausdorff matching [4, 36] and deformable parts models [21, 22, 26], that address this problem by allowing input images to undergo small spatial deformations before comparison. Before measuring the similarity of an
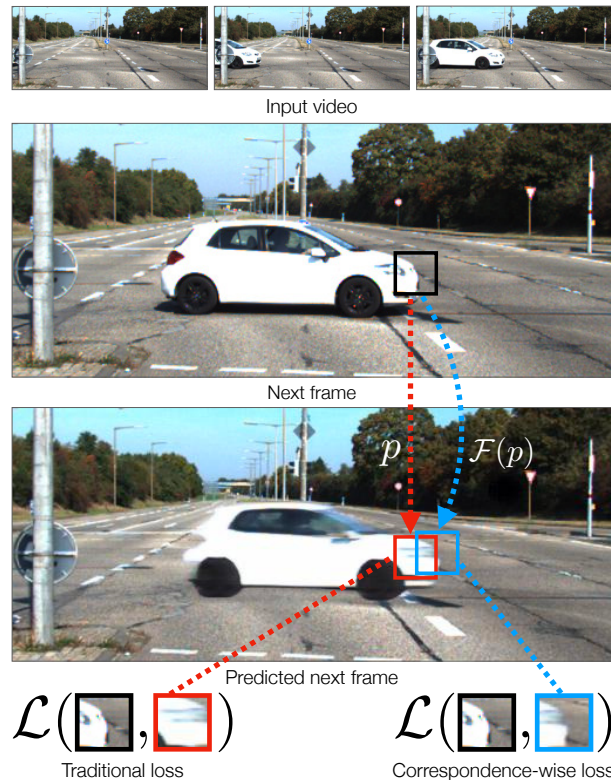

Figure 1. **Correspondence-wise losses.** We propose a similarity metric that provides robustness to small positional errors, and apply it to image generation. We put the predicted and ground truth images into correspondence via optical flow, then measure the similarity between each pixel p and its matching pixel $\mathcal{F}(p)$. Our metric leads to crisp predictions; it penalizes blurry, hedged images, like the one shown here, since they cannot be easily put into correspondence with the ground truth.

image and a template, these methods first geometrically align them, thereby obtaining robustness to small variations in position or shape.

Analogously, we propose a simple change to existing losses that makes them more robust to small positional errors. When comparing two images, we put them into correspondence using optical flow, then measure the similarity between *matching pairs of pixels*. Comparisons between the images therefore occur between pixel correspondences,

rather than pixels that reside in the same spatial positions (Figure 1), *i.e.* the loss is computed *correspondence-wise* rather than pixel-wise.

Despite its simplicity, our proposed "loss extension" leads to crisper and more perceptually accurate predictions. To obtain low loss, the predicted images must be easy to match with the ground truth via optical flow: every pixel in a target image requires a high-quality match in the predicted image, and vice versa. Blurry predictions tend to obtain high loss, since there is no simple, smooth flow field that puts them into correspondence with the ground truth. The loss also encourages objects to be placed in their correct positions, as positional mistakes lead to poor quality matches and occluded content, both of which incur penalties.

Since optical flow matching occurs within the loss function our approach does not require altering the design of the network itself. This is in contrast to popular *flow-based* video prediction architectures [27, 39, 58, 106] that produce a deformation field within the network, then generate an image by warping the input frames.

We demonstrate the effectiveness of our method in a number of ways:

- We show through experiments on a variety of video prediction tasks that our method significantly improves perceptual image quality. Our evaluation studies a variety of loss functions, including L_1 and L_2 distance and perceptual losses [28, 42]. These losses produce better results when paired with correspondence-wise prediction on egocentric driving datasets [13, 29].
- We obtain video prediction results that outperform a flow-based state-of-the-art video prediction method [95] on perceptual quality metrics for KITTI [29] and Cityscapes [13], despite using a simple, off-the-shelf network architecture.
- We apply our loss to the closely-related task of video interpolation [45], where we obtain significantly better results than an L_1 loss alone.
- We show that our method also improves performance of stochastic, variational autoencoder (VAE) video prediction architectures [15, 88].

2. Related Work

Video prediction. Early work in video prediction used recurrent networks to model long-range dependencies [44, 64, 65, 80, 83]. More recent work has focused on photorealistic video prediction using large convolutional networks. Lotter *et al.* [60] proposed a predictive coding method and applied it to driving videos. Wang *et al.* [92] predicted future semantic segmentation maps then translated them into images. Other work has improved image quality using adversarial losses [50, 52, 92], multiscale models [62], and recurrent networks with contextual aggregation [8]. These methods are complementary to ours, as our loss is archi-

tecture agnostic. Other work uses video prediction methods for model-based reinforcement learning [23, 30, 68]. Recently, Jayaraman *et al.* [40] proposed time-agnostic prediction, which gives a model flexibility to predict *any* of the future frames in a video. Our method proposes a similar mechanism, but in space, rather than time.

Another line of work has addressed the challenges of uncertainty in video prediction through stochastic models, such as variational autoencoders (VAEs) [48], that learn the full distribution of outcomes which is then sampled [1, 15, 100]. Notably, Denton and Fergus [15] introduced a recurrent variational autoencoder that used a learned prior distribution. This work was later extended by Villegas *et al.* [88], which introduced architectural changes and significantly increased the scale of the model, obtaining impressive results. Castrejón *et al.* [10] also observed that high capacity models improve generation results. Other work has introduced compositional models [102] and sparse predictions [31, 33, 91]. Our approach is complementary to this line of work; in our experiments we obtain benefits from using our loss in VAE-based video prediction [15, 88].

Flow-based video prediction. Rather than directly outputting images, many video prediction methods instead predict optical flow for each pixel and then synthesize a result by warping the input image. In early work, Patraucean *et al.* [72] predicted optical flow using a convolutional recurrent network. Liu *et al.* [58] predicted a 3D space-time flow field. Gao *et al.* [27] regressed motion using optical flow as pseudo-ground truth, inpainted occlusions, and used semantic segmentation. Later, Wu *et al.* [95] extended this approach, conditioning predictions on the trajectories of objects that their model segments and tracks. Recent work has used other motion representations, such as factorizing a scene into stationary and moving components [14, 87, 89], per-pixel kernels [24, 49, 67, 74, 90], or Eulerian motion [57]. Work in 3D view synthesis has adopted a similar approach, known as *appearance flow* [69, 69, 106]. Since these methods can only "copy and paste" existing content, they require special architectures to account for disocclusion and photometric changes. For example, state-of-the-art methods [27, 95] have motion estimation layers, internally perform warping via spatial transformers [39], and use separate inpainting modules [55, 103] to handle disocclusion. Since our approach only changes the loss function, it could in principle be combined with these architectures.

Perceptual losses. One way of reducing blur, commonly used in video prediction [11, 27, 95], is to use perceptual losses [28, 42]. These methods exploit the invariances learned for object recognition to provide robustness to small positional errors. However, because object recognition models learn only partial invariance, when positional errors are more than a few pixels they result in the same blurring artifacts seen in simpler, pixel-based losses. In our

experiments we show that this blurring can be reduced when applying our method.

Optical flow. Our method uses optical flow to obtain per-pixel correspondences. To solve this task, Lucas and Kanade [2, 61] make a brightness constancy assumption and solve a linearized model. Horn and Schunck [34] then proposed a smoothness prior for predicting flow. This approach was extended to use robust estimation methods [7, 20, 53, 81]. More recent methods use CNNs trained with supervised [25, 35, 37, 73, 82, 84, 101] or unsupervised learning [5, 43, 56, 75, 94, 104]. Teed and Deng [84] proposed an architecture for incrementally refining a flow field. While these works find space-time correspondences between frames, we instead use it to find correspondences between generated and ground truth images. In this way, our work is related to methods that use optical flow for other tasks, such as matching scenes [54], features [97], or objects [66, 77, 105].

Deformable matching. Our approach is inspired by classic work in image matching, particularly Chamfer [4] and Hausdorff [36] matching. These methods align a template to an image before comparison, providing robustness to small positional errors. A similar approach has been used in image retrieval [85], and part-based object detection [21, 22, 26, 107]. Single-image depth estimation has used analogous invariances to scale or space in its loss functions [18, 19]. Like these works, we allow images to undergo deformations before comparison, but we do so for *synthesis* instead of matching or detection.

Video frame interpolation. Frame interpolation shares many of the same challenges as video prediction, since the position and motion of objects are often uncertain. To tackle this problem, various models have been proposed that rely on optical flow [3, 41, 51, 70, 98], depth [3], or image kernels [12, 51]. Recently, Kalluri *et al.* [45] proposed a 3D CNN for interpolation. We augment this architecture with our loss and show improvements in performance.

3. Correspondence-wise Image Prediction

Our goal is to solve image prediction tasks where there is uncertainty in the positions of objects. To address this problem, we propose a "loss extension" that provides robustness to small positional misalignment. Given two images $\mathbf{x}$ and $\hat{\mathbf{x}}$, a traditional *pixel-wise* loss (*e.g.*, L_1 distance between images) can be written as:

$$\mathcal{L}_P(\mathbf{x}, \hat{\mathbf{x}}) = \frac{1}{|\mathcal{P}|} \sum_{p \in \mathcal{P}} \mathcal{L}(\mathbf{x}_p, \hat{\mathbf{x}}_p), \qquad (1)$$

where $\mathcal{L}$ is a *base* loss (e.g., L_1 for a pair of pixel intensities), $\mathbf{x}_p$ is the pixel color at location p in $\mathbf{x}$, and $\mathcal{P}$ is the set of all pixel indices.

Algorithm 1 Pseudocode in a PyTorch-like style for training an image prediction method with a correspondence-wise L_1 loss.

```
# Load a minibatch with source and target images
for (im_src, im_tgt) in loader:
 # Predict image using a network
 im_est = predict_image(im_src)

 # Estimate optical flow in both directions
 F_est = optical_flow(im_tgt, im_est)
 F_tgt = optical_flow(im_est, im_tgt)

 # Regularize the flow
 F_est = (1 - epsilon) * F_est
 F_tgt = (1 - epsilon) * F_tgt

 # Warp using bilinear filtering
 warp_est = warp(im_est, F_est)
 warp_tgt = warp(im_tgt, F_tgt)

 # Bidirectional loss with existing loss function
 loss1 = l1_loss(im_tgt, warp_est)
 loss2 = l1_loss(im_est, warp_tgt)
 loss = loss1 + loss2

 loss.backward()
```

`warp`: bilinear warping with an optical flow field.

In our approach, instead of comparing pixels at the same indices as in pixel-wise losses, we first compute pixel-to-pixel correspondences between the images, $\mathcal{F}(\mathbf{x}, \hat{\mathbf{x}})$, using optical flow. Then we compare each pixel $\mathbf{x}_p$ to its corresponding pixel $\hat{\mathbf{x}}_{F(p)}$, where $F(p)$ is the pixel in $\hat{\mathbf{x}}$ matched by the flow field. This loss can be written:

$$\mathcal{L}_C(\mathbf{x}, \hat{\mathbf{x}}) = \frac{1}{|\mathcal{P}|} \sum_{p \in \mathcal{P}} \mathcal{L}(\mathbf{x}_p, \hat{\mathbf{x}}_{F(p)}), \qquad (2)$$

We call the resulting loss a *correspondence-wise loss*. For example, when $\mathcal{L} = L_1$, we call it a correspondence-wise L_1 loss. The loss is illustrated in Figure 1.

We first detail the implementation of the loss in Section 3.1 and 3.2, and then investigate its properties in Section 3.3. Pseudocode for the full method, including all of the following implementation details, is provided in Alg. 1.

3.1. Regularization

Flow scaling. Models that minimize Equation 2 can fall into local optima during training, especially in the early stages when images are out-of-domain for the flow network. In addition, the warping process "snaps" objects to their exact ground truth positions, making it hard for the model to infer where to place objects in the generated image.

To address these issues, we introduce a small multiplicative decay to the flow field: $\mathcal{F}_R(\mathbf{x}_1, \mathbf{x}_2) = (1 - \epsilon)\mathcal{F}(\mathbf{x}_1, \mathbf{x}_2)$. This reduces long-range matching while encouraging the model to place objects closer to their true locations in the target image. In each training step, a model can decrease its loss if it moves an incorrectly placed object slightly closer to its true location. We use $\epsilon = 0.1$ and call this regularization strategy *flow scaling*.

Alternative methods. We also considered an alternative regularization strategy, inspired by Chamfer distance [4] and optical flow smoothness [43, 81] that directly penalizes

	20% Motion Range	30% Motion Range
L1		
Corr. L1		

(a) Average image under different losses

(b) Loss vs. Object Offset

Figure 2. **The effect of positional errors on different losses**. (a) We show images that minimize the expected error for a toy video prediction task. The car's location is sampled uniformly about the center at random (either 20% or 30% of the image width). The full range of the car's locations is indicated by the black bars. Note that the L_1 loss predictions are blurry, in particular for the 30% setting, whereas the *correspondence-wise* L_1 loss predictions are crisp in both cases. (b) We examine how the loss changes as a function of positional error (i.e., how far the predicted object is from its true location). Correspondence-wise losses increase smoothly with magnitude of the error.

the distance that each pixel moves in the flow field. When using this scheme, our regularization term is

$$\mathcal{L}_{reg}(\mathbf{u}) = \lambda_1 ||\mathbf{u}||^2 + \lambda_2 \mathcal{L}_{edge}(\nabla \mathbf{u}), \quad (3)$$

where $\mathcal{L}_{edge}$ is the edge-aware first-order smoothness penalty of Jonschkowski *et al.* [43], $\nabla \mathbf{u}$ are the flow gradients, and λ_i are weights. While we found this approach to be effective in some applications, we found that flow scaling generally performed better, and requires fewer hyperparameters (see Section 4.1.3).

3.2. Implementation details

Finding correspondences. In order to find $\mathcal{F}(\mathbf{x}, \hat{\mathbf{x}})$, we use RAFT [84], an optical flow network which predicts dense correspondences. In addition we have found that other models work, such as PWC-Net [82] (see Section 4.1.3).

Warp formulation. Equation 2 can be evaluated by iterating over pixel locations, calculating $\mathcal{F}(p)$, taking distances, and then averaging. However, in practice, we implement our method as an image warp, followed by a pixelwise loss:

$$\mathcal{L}_C(\mathbf{x}, \hat{\mathbf{x}}) \approx \mathcal{L}(\mathbf{x}, \text{warp}(\hat{\mathbf{x}}, F(\mathbf{x}, \hat{\mathbf{x}}))), \quad (4)$$

where $\text{warp}(\hat{\mathbf{x}}, \mathbf{u})$ is a backward warp of $\hat{\mathbf{x}}$ using the deformation field $\mathbf{u}$. Intuitively, the warp operation aligns pixels with their correspondences, after which we can apply an existing loss function. This formulation makes it straightforward to turn existing loss functions (e.g., perceptual losses [42] that operate on patches) into correspondence-wise losses, by warping and then applying the loss.

Symmetry. Following common practice for matching-based loss functions [36], we make the loss symmetric using $\mathcal{L}_{sym} = \mathcal{L}_C(\mathbf{x}, \hat{\mathbf{x}}) + \mathcal{L}_C(\hat{\mathbf{x}}, \mathbf{x})$. This discourages models

from generating superfluous content that has no correspondence in the target image.

3.3. Analyzing Correspondence-wise Prediction

To help understand how correspondence-wise losses address the challenges of positional uncertainty, we analyze its behavior on several simplified toy prediction tasks.

Motion uncertainty. We create a simple prediction task with inherent uncertainty in position. In the example shown in Figure 2, an object moves horizontally at unknown speed to a position uniformly distributed in a region near the image center. We ask what the optimal prediction is under various losses. For a loss $\mathcal{L}$, this is the image $\overline{\mathbf{x}}$ that minimizes the expected loss $\mathbb{E}_{\mathbf{x} \sim \mathcal{D}}[\mathcal{L}(\mathbf{x}, \overline{\mathbf{x}})]$, where $\mathcal{D}$ is the image distribution. We find $\overline{\mathbf{x}}$ using stochastic gradient descent.

For $\mathcal{L} = L_1$, the resulting image suffers from blurring[1]. By contrast, our correspondence-based L_1 loss leads to a sharp prediction, with the object located in the center of the distribution. More generally, our loss tends to favor *crisp predictions* that commit to a single position over blurry ones. This is for two reasons: (i) blurry predictions are harder to match than sharper images via a smooth flow field, and (ii) images far from their correct position are harder to match. The same behavior occurs for perceptual losses and on other scenes (see appendix).

Effect of positional error. Next, we asked how sensitive the loss is to position. For instance, could a lazy prediction method simply place objects in incorrect positions, nonetheless obtaining low loss when optical flow "fixes" the mistake, *e.g.*, a video prediction model that merely repeats the

---

[1]We note that there is an analytical solution for the L_1 loss, namely the median, which our SGD results obtain.

Figure 3. **Next-frame prediction with and without a correspondence-wise loss.** We show both $L_1 + \mathcal{L}_p$ and L_1 results, indicated by the left most column, on KITTI and Caltech. We highlight a noteworthy part of each result. Using a correspondence-wise loss results in sharper details (e.g., crisper road lines) and more robustness to large motions.

last frame? We visualize the loss as a function of an object's positional error in Figure 2, i.e. the loss incurred if the object were predicted a given offset from its true position. To reduce the effect of the background, we averaged the results over a large number of backgrounds (see appendix for details). We see that our loss, in fact, *steadily increases with positional offset*. Moreover, the global minimum remains the same. Three reasons we found for this are: (i) Almost any incorrect prediction will have *occluded or extraneous content* that incurs a large loss. (ii) When flow regularization is used, the model is explicitly penalized for incorrect predictions. (iii) The implicit smoothness prior in optical flow estimation trades off reconstruction error for small, simple motions; flow methods will tend to choose matches that incur large reconstruction error over those with large flow values.

4. Results

Our goal is to understand how correspondence-wise losses differ from their pixelwise counterparts, and to evaluate their effectiveness. To do this we perform experiments on video prediction, as well as frame interpolation. We ablate pixelwise and correspondence-wise losses across various tasks, datasets, architectures, and metrics. In addition, we compare models trained with our loss to the state-of-the-art methods.

Base Loss	Corr.	KITTI			Cityscapes		Caltech	
		SSIM ↑	LPIPS ↓	2AFC ↑	SSIM ↑	LPIPS ↓	SSIM ↑	LPIPS ↓
L_1	-	0.563	0.438	13.52	**0.820**	0.231	0.733	0.250
L_1	✓	**0.586**	**0.359**	**14.25**	0.819	**0.198**	**0.734**	**0.216**
L_2^2	-	0.544	0.499	13.36	0.801	0.291	0.701	0.330
L_2^2	✓	**0.563**	**0.403**	**15.16**	**0.812**	**0.212**	**0.707**	**0.249**
$L_1 + \mathcal{L}_p$	-	0.545	0.213	14.46	**0.816**	0.092	**0.717**	**0.139**
$L_1 + \mathcal{L}_p$	✓	**0.548**	**0.191**	**20.19**	0.810	**0.090**	0.702	0.141

Table 1. **Pixel-wise vs. correspondence-wise losses.** We compare correspondence-wise losses to their pixelwise counter parts, showing its efficacy on KITTI, Cityscapes, and Caltech. The second column indicates the usage of a correspondence-wise loss. 2AFC is the rate at which humans chose the generated video over the true video in a real-or-fake study.

4.1. Video Prediction

Models. We consider both deterministic and stochastic video prediction architectures. For simplicity, our deterministic model is based on the widely-used residual network [32, 42] from Wang *et al.* [93], except we replace 2D convolutions by 3D convolutions [9, 86] to process temporal information, and we replace transposed convolutions with upsampling followed by 2D convolutions to avoid checkerboard artifacts (see appendix for full network architecture). Multi-frame prediction is performed by recursively feeding output images back into the model, as in [27, 60, 95].

The stochastic model we use is the **SVG** model intro-

duced by Denton and Fergus [15], a CNN with a learned prior and LSTM layers. In addition, we adopt extensions to the SVG model proposed by Villegas *et al.* [88], resulting in a model we call **SVG++**. This is a large-scale model that offers architectural improvements over SVG, and which obtains strong performance on the KITTI dataset. Following Villegas *et al.* we modify SVG by adding convolutional LSTM layers [78], although we keep the L_2 loss from [15] (see appendix for details). Since our goal is to understand the influence of a correspondence-wise loss, rather than to obtain state-of-the-art performance, we base our model on a medium-sized variant with hyperparameters $K = 2$ and $M = 2$ [88], which obtains strong performance yet is trainable on ordinary multi-GPU computing infrastructure.

For optical flow estimation, we use RAFT [84], with the publicly available checkpoint that was trained on Flying Chairs [17] and FlyingThings [63]. While our results may be improved by using a version trained on driving videos from KITTI [29], we choose not to do so in our experiments to avoid the use of domain-specific supervision.

Losses. We use three base losses: **1)** the L_1 loss, **2)** the MSE loss (L_2^2), and **3)** the L_1 loss in equal weight with a perceptual loss, $\mathcal{L}_p$, ($L_1 + \mathcal{L}_p$) that uses VGG-19 pretrained with ImageNet features [42, 76, 79].

Metrics. To evaluate predictions, we use SSIM, LPIPS, and a two alternative forced choice human study (see appendix for details).

Datasets. We evaluate on three standard video prediction datasets: KITTI [29], Caltech Pedestrian [16], and Cityscapes [13] (see appendix for details).

4.1.1 Pixelwise vs. Correspondence-wise losses

To understand the effect of our extension of pixelwise losses, we train video prediction models to predict three future frames from three previous frames with both losses, using three separate base losses: L_1, L_2^2 (MSE), and $L_1 + \mathcal{L}_p$. We evaluate at a resolution of 512×256 on the KITTI dataset, 384×288 on Caltech, and 512×256 of Cityscapes, all with three frames of input. For consistency with the KITTI experiments, we sample the Caltech dataset at 10 Hz, as in Lotter *et al.* [60]. In addition, for the $L_1 + \mathcal{L}_p$ loss, we warm start with a pixelwise loss for an epoch, which significantly improves convergence.

The results can be found in Table 1. The correspondence-wise variants of the loss outperform their pixelwise counterparts on almost all metrics, datasets, and losses. In addition, we see significant improvements in the qualitative results across all scenarios, as seen in Figure 3.

4.1.2 Comparison to State-of-the-Art Methods

We follow the evaluation protocol of Wu *et al.* [95], using the KITTI and Cityscapes datasets, and compare with a number of recent video prediction methods:

Voxel Flow [58], a motion-synthesis method based on 3D space-time flow, **MCnet** [89], a convolutional LSTM model that decomposes stationary and moving components, **Vid2Vid** [92], a two-stage method that first synthesizes semantic masks and then translates the masks to real images, and **OMP** [95], a state-of-the-art method that combines copy-and-paste prediction with inpainting, object tracking, occlusion estimation, and adversarial training. We also show results with **PredNet** [60], a convolutional pixel-based architecture inspired by predictive coding. We note that these architectures are specialized to the video prediction task, and can be relatively complex. For example, OMP uses off-the-shelf instance segmentation and semantic segmentation networks [96, 108], an inpainting network [103], and a background-prediction network. It also takes optical flow as input [82], tracks objects, and uses adversarial training. By contrast, we are interested in seeing how well a simple image prediction network can do when using our loss.

Following Wu *et al.* [95], we use 832×256 images on KITTI and 1024×512 on Cityscapes, using the same train-test split and data augmentation, and condition our models on four frames of an input video and predict five future frames on KITTI and 10 on Cityscapes. We use our ResNet variant with an $L_1 + \mathcal{L}_p$ correspondence-wise loss.

Despite our method's simplicity, we found that it significantly outperformed previous methods with complex flow-based architectures on both metrics (Table 2). Our simple architecture, trained with a correspondence-wise loss, obtained higher scores on both SSIM and LPIPS consistently across all time steps.

In Figure 4, we show three informative qualitative results generated by our model, as compared to frames generated by OMP. We have highlighted challenging regions in each video. In contrast to our model, OMP works by deforming an input image using a predicted optical flow field, resulting in warping artifacts when there are errors. Here, OMP suffers when there are small objects undergoing large motions (*e.g.*, the poles in the first video) and irregular geometry (*e.g.*, second video). Interestingly, both models produce errors in disoccluded regions (*e.g.*, car in first video).

4.1.3 Additional Ablations

We present additional ablations on network architecture, regularization scheme, and flow method for the video prediction task.

Stochastic models. We demonstrate the versatility of our loss by evaluating its performance with stochastic video prediction architectures on KITTI in Table 3. Evaluation metrics on both the SVG and SVG++ significantly improve when we use a correspondence-wise reconstruction loss.

Effect of regularization. In Section 3.1 we introduced two methods of regularization. Table 4 evaluates these two
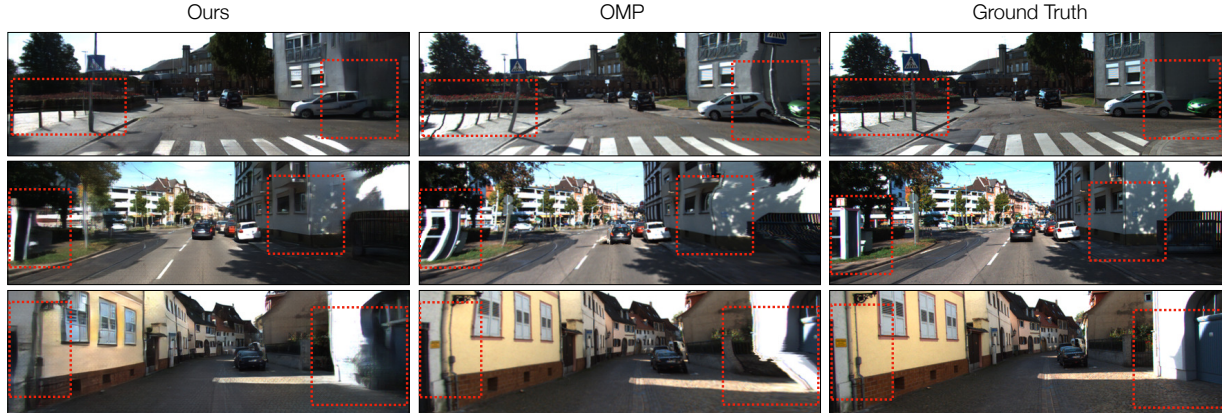
Figure 4. **Comparison with state-of-the-art multi-frame video prediction.** We show the result of our method and OMP [95] on predicting frame $t = 5$. Areas marked with a box are challenging examples.

| | KITTI | | | | | | Cityscapes | | | | | |
| | Next Frame | | Next 3 Frames | | Next 5 Frames | | Next Frame | | Next 5 Frames | | Next 10 Frames | |
Model	SSIM ↑	LPIPS↓	SSIM ↑	LPIPS↓	SSIM ↑	LPIPS↓	SSIM ↑	LPIPS↓	SSIM ↑	LPIPS↓	SSIM ↑	LPIPS↓
PredNet [60]	0.563	0.553	0.514	0.586	0.475	0.629	0.840	0.260	0.752	0.360	0.663	0.522
MCNET [89]	0.753	0.240	0.635	0.317	0.554	0.373	0.897	0.189	0.706	0.373	0.597	0.451
Voxel Flow [58]	0.539	0.324	0.469	0.374	0.426	0.415	0.839	0.174	0.711	0.288	0.634	0.366
Vid2Vid [92]	-	-	-	-	-	-	0.882	0.106	0.751	0.201	0.669	0.271
OMP [95]	0.792	0.185	0.676	0.246	0.607	0.304	0.891	**0.085**	0.757	0.165	0.674	0.233
Ours	**0.820**	**0.172**	**0.730**	**0.220**	**0.667**	**0.259**	**0.928**	**0.085**	**0.839**	**0.150**	**0.751**	**0.217**

Table 2. **Comparisons with other methods.** We compare with state-of-the-art methods on multi-frame KITTI and Cityscapes video prediction, as well PredNet, an older non-flow based method. Note that in keeping with previous works, we present the *average* of the first N frames and we use the linear variant of LPIPS. We take baseline numbers directly from Wu *et al.* [95].

| | | Next Frame | | Next 5 Frames | | Next 10 Frames | |
Model	Corr.	SSIM ↑	LPIPS ↓	SSIM ↑	LPIPS ↓	SSIM ↑	LPIPS ↓
SVG [15]	-	0.389	0.478	0.342	0.509	0.312	0.527
	✓	**0.400**	**0.389**	**0.348**	**0.382**	**0.313**	**0.386**
SVG++ [88]	-	0.848	0.079	0.626	0.196	0.489	0.287
	✓	**0.849**	**0.072**	**0.628**	**0.186**	**0.490**	**0.276**

Table 3. **Stochastic model evaluation.** We show results with and without our correspondence-wise loss ("Corr.") while using stochastic video prediction architectures.

| | | KITTI | | Caltech | | Cityscapes | |
Base Loss	Reg	SSIM ↑	LPIPS ↓	SSIM ↑	LPIPS ↓	SSIM ↑	LPIPS ↓
L_1	Scale	**0.586**	**0.359**	**0.734**	**0.216**	0.819	**0.198**
L_1	Mag.	0.573	0.407	0.723	0.249	**0.826**	0.227
$L_1 + \mathcal{L}_p$	Scale	**0.548**	**0.191**	0.702	**0.141**	0.810	**0.090**
$L_1 + \mathcal{L}_p$	Mag.	**0.548**	0.206	**0.707**	0.149	**0.823**	0.098

Table 4. **Regularization for correspondence-wise losses.** We ablate the flow scaling and flow magnitude penalty regularization methods, "Scale" and "Mag." in the table respectively. These experiments correspond to the ones in Table 1.

Metric	Pixelwise	Affine	Hom.	PWC-net	RAFT
LPIPS ↓	0.438	0.393	0.393	0.364	**0.359**
SSIM ↑	0.563	0.568	0.569	**0.586**	**0.586**

Table 5. **Influence of flow method.** We show image quality as a function of flow method with an L_1 base loss on KITTI (the setup in the first two rows of Table 1).

was challenging, and that it often had poor convergence, which we addressed by training with a warm start from a base loss for 5 epochs.

Quality of flow method. In Table 5 we test the influence of the quality of the flow method by comparing RAFT [84] and PWC-Net [82]. In addition, we simulated less powerful flow methods by simply fitting a homography or an affine transformation to the predicted RAFT flow, and using these fitted transforms as our flow. We found that using a correspondence-wise loss with better quality flow methods produces better downstream image generation performance.

4.2. Video Frame Interpolation

To further evaluate our loss's ability to handle spatial uncertainty, we conduct frame interpolation experiments. Given two frames of context, the goal is to predict the intermediate frame. We use the FLAVR architecture [45] with both an L_1 loss and an L_1 correspondence-wise loss and

approaches. The *magnitude penalty* approach outperforms the *scaling* approach in some circumstances. However, we found that having to tune each of its weighting parameters
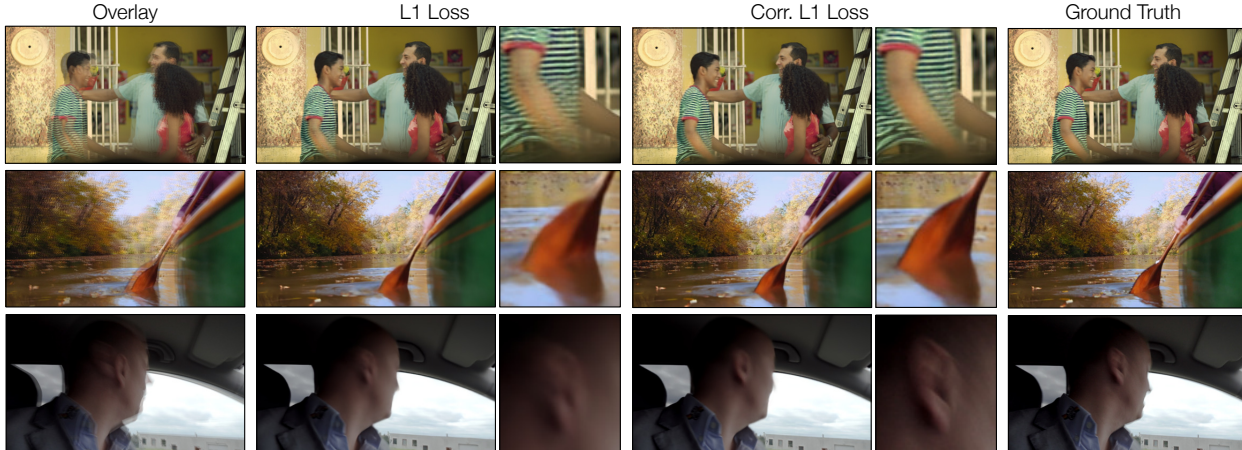
Figure 5. **Frame interpolation qualitative results.** We visualize frame interpolation results from Vimeo-90K for a FLAVR model trained on an L_1 loss and a correspondence-wise L_1 loss. The leftmost column shows the two context frames, overlaid on top of each other. We highlight a noteworthy region of each result.

Method	Frames	LPIPS ($\downarrow$)	PSNR ($\uparrow$)	SSIM ($\uparrow$)
DVF [59]	2	-	27.27	0.893
SuperSloMo [41]	2	-	32.90	0.957
SepConv [67]	2	-	33.60	0.944
CAIN [12]	2	-	33.93	0.964
BMBC [70]	2	-	34.76	0.965
AdaCoF [51]	2	-	35.40	0.971
FLAVR [45]	4	0.0248	36.30	0.975
FLAVR [45]	2	0.0297	34.96	0.970
FLAVR + Ours	2	0.0268	35.13	0.970

Table 6. **Frame interpolation evaluation.** We replace the L_1 loss in FLAVR with a correspondence-wise L_1 loss. We report baseline PSNR and SSIM numbers from Kalluri *et al.* [45].

we evaluate using SSIM, PSNR, and LPIPS on the Vimeo-90K septuplet dataset [99]. These results, along with various two-frame baselines and a FLAVR model trained on four frames, are presented in Table 6.

We trained our model using the publicly released code from FLAVR, replacing their loss with our correspondence-wise loss. To be more consistent with the baselines and to study a scenario with more uncertainty, we use two context frames instead of four. Specifically we use the two middle frames, closest to the ground truth, in the septuplets. We train for 120 epochs, use a batch size of 32, and learning rate of 0.0002 with Adam [47]. These hyperparameters are the same for all FLAVR models.

The FLAVR model trained with our correspondence-wise L_1 loss outperforms the model with only L_1 loss on LPIPS and PSNR, ties on SSIM, and leads to crisper interpolations (Fig. 5). In addition, the model outperforms all but one baseline and is even competitive with the FLAVR model trained on four frames.

5. Discussion

Limitations. One drawback of our method is that it adds overhead at training time, since it requires evaluating an optical flow network twice per example. In our ablation experiments, corresponding to Table 1, we found an increase of approximately 30%-50% wall time. However, test time inference speed is unaffected as our method is a loss function. Additionally, our method requires the optical flow model to successfully match the predicted and ground truth images, which may not be possible in some tasks.

Conclusion. In this paper we have proposed a *correspondence-wise* loss for image generation, and applied it to the task of video prediction and frame interpolation. We show through extensive ablations that our loss improves image quality over pixelwise losses on various metrics, architectures, datasets, and tasks. Despite the simplicity of our approach, it outperforms several recent, highly engineered video prediction methods when used with simple off-the-shelf architectures.

We see our work opening two directions. The first is to unify "copy-and-paste" synthesis methods [72, 106] and correspondence-wise methods by designing architectures that can obtain the benefits of both techniques. The second is to use correspondence-wise image prediction in other tasks where positional uncertainty currently leads to blurry results, such as viewpoint synthesis and image translation.

References

[1] Mohammad Babaeizadeh, Chelsea Finn, Dumitru Erhan, Roy H. Campbell, and Sergey Levine. Stochastic variational video prediction, 2017. 2

[2] Simon Baker and Iain Matthews. Lucas-kanade 20 years on: A unifying framework. In *International Journal of Computer Vision*, pages 221–255, Feb. 2004. 3

[3] Wenbo Bao, Wei-Sheng Lai, Chao Ma, Xiaoyun Zhang, Zhiyong Gao, and Ming-Hsuan Yang. Depth-aware video frame interpolation, 2019. 3

[4] HG Barrow, JM Tenenbaum, RC Bolles, and HCf Wolf. Parametric correspondence and chamfer matching: Two new techniques for image matching. In *Proceedings: Image Understanding Workshop*, pages 21–27. Science Applications, Inc Arlington, VA, 1977. 1, 3

[5] Zhangxing Bian, Allan Jabri, Alexei A. Efros, and Andrew Owens. Learning pixel trajectories with multiscale contrastive random walks. *CVPR*, 2022. 3

[6] Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale gan training for high fidelity natural image synthesis, 2018. 1

[7] Thomas Brox, Andrés Bruhn, Nils Papenberg, and Joachim Weickert. High accuracy optical flow estimation based on a theory for warping. In *European conference on computer vision*, pages 25–36. Springer, 2004. 3

[8] Wonmin Byeon, Qin Wang, Rupesh Kumar Srivastava, and Petros Koumoutsakos. Contextvp: Fully context-aware video prediction. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 753–769, 2018. 2

[9] Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset, 2017. 5

[10] Lluis Castrejon, Nicolas Ballas, and Aaron Courville. Improved conditional vrnns for video prediction, 2019. 2

[11] Qifeng Chen and Vladlen Koltun. Photographic image synthesis with cascaded refinement networks, 2017. 1, 2

[12] Myungsub Choi, Heewon Kim, Bohyung Han, Ning Xu, and Kyoung Mu Lee. Channel attention is all you need for video frame interpolation. In *AAAI*, 2020. 3, 8

[13] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset for semantic urban scene understanding. In *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 2, 6

[14] Emily Denton and Vighnesh Birodkar. Unsupervised learning of disentangled representations from video, 2017. 2

[15] Emily Denton and Rob Fergus. Stochastic video generation with a learned prior. *arXiv preprint arXiv:1802.07687*, 2018. 2, 6, 7

[16] P. Dollár, C. Wojek, B. Schiele, and P. Perona. Pedestrian detection: A benchmark. In *CVPR*, June 2009. 6

[17] A. Dosovitskiy, P. Fischer, E. Ilg, P. Häusser, C. Hazırbaş, V. Golkov, P. v.d. Smagt, D. Cremers, and T. Brox. Flownet: Learning optical flow with convolutional networks. In *IEEE International Conference on Computer Vision (ICCV)*, 2015. 6

[18] David Eigen, Christian Puhrsch, and Rob Fergus. Depth map prediction from a single image using a multi-scale deep network. *arXiv preprint arXiv:1406.2283*, 2014. 3

[19] Haoqiang Fan, Hao Su, and Leonidas J Guibas. A point set generation network for 3d object reconstruction from a single image. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 605–613, 2017. 3

[20] PF Felzenszwalb and DR Huttenlocher. Efficient belief propagation for early vision. In *Proceedings of the 2004 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 2004. CVPR 2004.*, volume 1, pages I–I. IEEE, 2004. 3

[21] Pedro F Felzenszwalb, Ross B Girshick, David McAllester, and Deva Ramanan. Object detection with discriminatively trained part-based models. *IEEE transactions on pattern analysis and machine intelligence*, 32(9):1627–1645, 2009. 1, 3

[22] Pedro F Felzenszwalb and Daniel P Huttenlocher. Pictorial structures for object recognition. *International journal of computer vision*, 61(1):55–79, 2005. 1, 3

[23] Chelsea Finn, Ian Goodfellow, and Sergey Levine. Unsupervised learning for physical interaction through video prediction, 2016. 2

[24] Chelsea Finn, Ian Goodfellow, and Sergey Levine. Unsupervised learning for physical interaction through video prediction. *arXiv preprint arXiv:1605.07157*, 2016. 2

[25] Philipp Fischer, Alexey Dosovitskiy, Eddy Ilg, Philip Häusser, Caner Hazırbaş, Vladimir Golkov, Patrick Van der Smagt, Daniel Cremers, and Thomas Brox. Flownet: Learning optical flow with convolutional networks. *arXiv preprint arXiv:1504.06852*, 2015. 3

[26] Martin A Fischler and Robert A Elschlager. The representation and matching of pictorial structures. *IEEE Transactions on computers*, 100(1):67–92, 1973. 1, 3

[27] Hang Gao, Huazhe Xu, Qi-Zhi Cai, Ruth Wang, Fisher Yu, and Trevor Darrell. Disentangling propagation and generation for video prediction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9006–9015, 2019. 2, 5

[28] Leon A Gatys, Alexander S Ecker, and Matthias Bethge. Image style transfer using convolutional neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2414–2423, 2016. 2

[29] Andreas Geiger, Philip Lenz, Christoph Stiller, and Raquel Urtasun. Vision meets robotics: The kitti dataset. *International Journal of Robotics Research (IJRR)*, 2013. 2, 6

[30] David Ha and Jürgen Schmidhuber. Recurrent world models facilitate policy evolution. In *Advances in Neural Information Processing Systems 31*, pages 2451–2463. Curran Associates, Inc., 2018. https://worldmodels.github.io. 2

[31] Zekun Hao, Xun Huang, and Serge Belongie. Controllable video generation with sparse trajectories. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7854–7863, 2018. 2

[32] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 5

[33] Yung-Han Ho, Chuan-Yuan Cho, Wen-Hsiao Peng, and Guo-Lun Jin. Sme-net: Sparse motion estimation for parametric video prediction through reinforcement learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10462–10470, 2019. 2

[34] Berthold Horn and Brain Schunck. Determining optical flow. In *Artificial Intelligence*, pages 185–203, 1981. 3

[35] Tak-Wai Hui, Xiaoou Tang, and Chen Change Loy. Liteflownet: A lightweight convolutional neural network for optical flow estimation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 8981–8989, 2018. 3

[36] Daniel P Huttenlocher, Gregory A. Klanderman, and William J Rucklidge. Comparing images using the hausdorff distance. *IEEE Transactions on pattern analysis and machine intelligence*, 15(9):850–863, 1993. 1, 3, 4

[37] Eddy Ilg, Nikolaus Mayer, Tonmoy Saikia, Margret Keuper, Alexey Dosovitskiy, and Thomas Brox. Flownet 2.0: Evolution of optical flow estimation with deep networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2462–2470, 2017. 3

[38] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A. Efros. Image-to-image translation with conditional adversarial networks, 2016. 1

[39] Max Jaderberg, Karen Simonyan, Andrew Zisserman, and Koray Kavukcuoglu. Spatial transformer networks, 2015. 2

[40] Dinesh Jayaraman, Frederik Ebert, Alexei A. Efros, and Sergey Levine. Time-agnostic prediction: Predicting predictable video frames, 2018. 2

[41] Huaizu Jiang, Deqing Sun, Varun Jampani, Ming-Hsuan Yang, Erik Learned-Miller, and Jan Kautz. Super slomo: High quality estimation of multiple intermediate frames for video interpolation, 2018. 3, 8

[42] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution, 2016. 2, 4, 5, 6

[43] Rico Jonschkowski, Austin Stone, Jonathan T Barron, Ariel Gordon, Kurt Konolige, and Anelia Angelova. What matters in unsupervised optical flow. *arXiv preprint arXiv:2006.04902*, 2020. 3, 4

[44] Nal Kalchbrenner, Aäron Oord, Karen Simonyan, Ivo Danihelka, Oriol Vinyals, Alex Graves, and Koray Kavukcuoglu. Video pixel networks. In *International Conference on Machine Learning*, pages 1771–1779. PMLR, 2017. 2

[45] Tarun Kalluri, Deepak Pathak, Manmohan Chandraker, and Du Tran. Flavr: Flow-agnostic video representations for fast frame interpolation, 2021. 2, 3, 7, 8

[46] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks, 2018. 1

[47] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 8

[48] Diederik P Kingma and Max Welling. Auto-encoding variational bayes, 2013. 2

[49] Shu Kong and Charless Fowlkes. Multigrid predictive filter flow for unsupervised learning on videos. *arXiv preprint arXiv:1904.01693*, 2019. 2

[50] Alex X. Lee, Richard Zhang, Frederik Ebert, Pieter Abbeel, Chelsea Finn, and Sergey Levine. Stochastic adversarial video prediction, 2018. 2

[51] Hyeongmin Lee, Taeoh Kim, Tae young Chung, Daehyun Pak, Yuseok Ban, and Sangyoun Lee. Adacof: Adaptive collaboration of flows for video frame interpolation, 2020. 3, 8

[52] Xiaodan Liang, Lisa Lee, Wei Dai, and Eric P Xing. Dual motion gan for future-flow embedded video prediction. In *proceedings of the IEEE international conference on computer vision*, pages 1744–1752, 2017. 2

[53] Ce Liu et al. *Beyond pixels: exploring new representations and applications for motion analysis*. PhD thesis, Massachusetts Institute of Technology, 2009. 3

[54] Ce Liu, Jenny Yuen, and Antonio Torralba. Sift flow: Dense correspondence across scenes and its applications. *IEEE transactions on pattern analysis and machine intelligence*, 33(5):978–994, 2010. 3

[55] Guilin Liu, Fitsum A Reda, Kevin J Shih, Ting-Chun Wang, Andrew Tao, and Bryan Catanzaro. Image inpainting for irregular holes using partial convolutions. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 85–100, 2018. 2

[56] Pengpeng Liu, Irwin King, Michael R Lyu, and Jia Xu. Ddflow: Learning optical flow with unlabeled data distillation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 8770–8777, 2019. 3

[57] Wenqian Liu, Abhishek Sharma, Octavia Camps, and Mario Sznaier. Dyan: A dynamical atoms-based network for video prediction. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 170–185, 2018. 2

[58] Ziwei Liu, Raymond A Yeh, Xiaoou Tang, Yiming Liu, and Aseem Agarwala. Video frame synthesis using deep voxel flow. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 4463–4471, 2017. 2, 6, 7

[59] Ziwei Liu, Raymond A. Yeh, Xiaoou Tang, Yiming Liu, and Aseem Agarwala. Video frame synthesis using deep voxel flow, 2017. 8

[60] William Lotter, Gabriel Kreiman, and David Cox. Deep predictive coding networks for video prediction and unsupervised learning, 2016. 2, 5, 6, 7

[61] Bruce D Lucas, Takeo Kanade, et al. An iterative image registration technique with an application to stereo vision. Vancouver, British Columbia, 1981. 3

[62] Michael Mathieu, Camille Couprie, and Yann LeCun. Deep multi-scale video prediction beyond mean square error. *arXiv preprint arXiv:1511.05440*, 2015. 2

[63] N. Mayer, E. Ilg, P. Häusser, P. Fischer, D. Cremers, A. Dosovitskiy, and T. Brox. A large dataset to train convolutional networks for disparity, optical flow, and scene flow estimation. In *IEEE International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. arXiv:1512.02134. 6

[64] Vincent Michalski, Roland Memisevic, and Kishore Konda. Modeling deep temporal dependencies with recurrent grammar cells"". In *Advances in neural information processing systems*, pages 1925–1933, 2014. 2

[65] Roni Mittelman, Benjamin Kuipers, Silvio Savarese, and Honglak Lee. Structured recurrent temporal restricted boltzmann machines. In *International Conference on Machine Learning*, pages 1647–1655, 2014. 2

[66] Hossein Mobahi, Ce Liu, and William T Freeman. A compositional model for low-dimensional image set representation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1322–1329, 2014. 3

[67] Simon Niklaus, Long Mai, and Feng Liu. Video frame interpolation via adaptive convolution. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 670–679, 2017. 2, 8

[68] Junhyuk Oh, Xiaoxiao Guo, Honglak Lee, Richard Lewis, and Satinder Singh. Action-conditional video prediction using deep networks in atari games, 2015. 2

[69] Eunbyung Park, Jimei Yang, Ersin Yumer, Duygu Ceylan, and Alexander C Berg. Transformation-grounded image generation network for novel 3d view synthesis. In *Proceedings of the ieee conference on computer vision and pattern recognition*, pages 3500–3509, 2017. 2

[70] Junheum Park, Keunsoo Ko, Chul Lee, and Chang-Su Kim. Bmbc:bilateral motion estimation with bilateral cost volume for video interpolation, 2020. 3, 8

[71] Taesung Park, Ming-Yu Liu, Ting-Chun Wang, and Jun-Yan Zhu. Semantic image synthesis with spatially-adaptive normalization. 2019. 1

[72] Viorica Patraucean, Ankur Handa, and Roberto Cipolla. Spatio-temporal video autoencoder with differentiable memory. *arXiv preprint arXiv:1511.06309*, 2015. 2, 8

[73] Anurag Ranjan and Michael J Black. Optical flow estimation using a spatial pyramid network. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4161–4170, 2017. 3

[74] Fitsum A Reda, Guilin Liu, Kevin J Shih, Robert Kirby, Jon Barker, David Tarjan, Andrew Tao, and Bryan Catanzaro. Sdc-net: Video prediction using spatially-displaced convolution. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 718–733, 2018. 2

[75] Zhe Ren, Junchi Yan, Bingbing Ni, Bin Liu, Xiaokang Yang, and Hongyuan Zha. Unsupervised deep learning for optical flow estimation. In *Thirty-First AAAI Conference on Artificial Intelligence*, 2017. 3

[76] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *International journal of computer vision*, 115(3):211–252, 2015. 6

[77] Xi Shen, François Darmon, Alexei A Efros, and Mathieu Aubry. Ransac-flow: generic two-stage image alignment. *arXiv preprint arXiv:2004.01526*, 2020. 3

[78] Xingjian Shi, Zhourong Chen, Hao Wang, Dit-Yan Yeung, Wai-Kin Wong, and Wang-chun Woo. Convolutional lstm network: A machine learning approach for precipitation nowcasting. *Advances in neural information processing systems*, 28:802–810, 2015. 6

[79] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. 6

[80] 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

[81] Deqing Sun, Stefan Roth, and Michael J Black. Secrets of optical flow estimation and their principles. In *2010 IEEE computer society conference on computer vision and pattern recognition*, pages 2432–2439. IEEE, 2010. 3

[82] Deqing Sun, Xiaodong Yang, Ming-Yu Liu, and Jan Kautz. Pwc-net: Cnns for optical flow using pyramid, warping, and cost volume, 2017. 3, 4, 6, 7

[83] Ilya Sutskever, Geoffrey E Hinton, and Graham W Taylor. The recurrent temporal restricted boltzmann machine. In *Advances in neural information processing systems*, pages 1601–1608, 2009. 2

[84] Zachary Teed and Jia Deng. Raft: Recurrent all-pairs field transforms for optical flow, 2020. 3, 4, 6, 7

[85] Antonio Torralba, Rob Fergus, and William T Freeman. 80 million tiny images: A large data set for nonparametric object and scene recognition. *IEEE transactions on pattern analysis and machine intelligence*, 30(11):1958–1970, 2008. 3

[86] Du Tran, Heng Wang, Lorenzo Torresani, Jamie Ray, Yann LeCun, and Manohar Paluri. A closer look at spatiotemporal convolutions for action recognition, 2017. 5

[87] Sergey Tulyakov, Ming-Yu Liu, Xiaodong Yang, and Jan Kautz. Mocogan: Decomposing motion and content for video generation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1526–1535, 2018. 2

[88] Ruben Villegas, Arkanath Pathak, Harini Kannan, Dumitru Erhan, Quoc V. Le, and Honglak Lee. High fidelity video prediction with large stochastic recurrent neural networks, 2019. 2, 6, 7

[89] Ruben Villegas, Jimei Yang, Seunghoon Hong, Xunyu Lin, and Honglak Lee. Decomposing motion and content for natural video sequence prediction, 2017. 2, 6, 7

[90] Carl Vondrick and Antonio Torralba. Generating the future with adversarial transformers. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1020–1028, 2017. 2

[91] Jacob Walker, Carl Doersch, Abhinav Gupta, and Martial Hebert. An uncertain future: Forecasting from static images using variational autoencoders. In *European Conference on Computer Vision*, pages 835–851. Springer, 2016. 2

[92] Ting-Chun Wang, Ming-Yu Liu, Jun-Yan Zhu, Guilin Liu, Andrew Tao, Jan Kautz, and Bryan Catanzaro. Video-to-video synthesis. *arXiv preprint arXiv:1808.06601*, 2018. 2, 6, 7

[93] 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, 2017. 5

[94] Yang Wang, Yi Yang, Zhenheng Yang, Liang Zhao, Peng Wang, and Wei Xu. Occlusion aware unsupervised learning of optical flow. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4884–4893, 2018. 3

[95] Yue Wu, Rongrong Gao, Jaesik Park, and Qifeng Chen. Future video synthesis with object motion prediction. In *CVPR*, 2020. 2, 5, 6, 7

[96] Yuwen Xiong, Renjie Liao, Hengshuang Zhao, Rui Hu, Min Bai, Ersin Yumer, and Raquel Urtasun. Upsnet: A unified panoptic segmentation network, 2019. 6

[97] Yuwen Xiong, Mengye Ren, Wenyuan Zeng, and Raquel Urtasun. Self-supervised representation learning from flow equivariance, 2021. 3

[98] Xiangyu Xu, Li Siyao, Wenxiu Sun, Qian Yin, and Ming-Hsuan Yang. Quadratic video interpolation, 2019. 3

[99] Tianfan Xue, Baian Chen, Jiajun Wu, Donglai Wei, and William T Freeman. Video enhancement with task-oriented flow. *International Journal of Computer Vision (IJCV)*, 127(8):1106–1125, 2019. 8

[100] Tianfan Xue, Jiajun Wu, Katherine Bouman, and Bill Freeman. Visual dynamics: Probabilistic future frame synthesis via cross convolutional networks. In *Advances in neural information processing systems*, pages 91–99, 2016. 2

[101] Gengshan Yang and Deva Ramanan. Volumetric correspondence networks for optical flow. In *Advances in neural information processing systems*, pages 794–805, 2019. 3

[102] Yufei Ye, Maneesh Singh, Abhinav Gupta, and Shubham Tulsiani. Compositional video prediction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10353–10362, 2019. 2

[103] Jiahui Yu, Zhe Lin, Jimei Yang, Xiaohui Shen, Xin Lu, and Thomas S. Huang. Generative image inpainting with contextual attention, 2018. 2, 6

[104] Jason J. Yu, Adam W. Harley, and Konstantinos G. Derpanis. Back to basics: Unsupervised learning of optical flow via brightness constancy and motion smoothness. In *Computer Vision - ECCV 2016 Workshops, Part 3*, 2016. 3

[105] Tinghui Zhou, Yong Jae Lee, Stella X Yu, and Alyosha A Efros. Flowweb: Joint image set alignment by weaving consistent, pixel-wise correspondences. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1191–1200, 2015. 3

[106] Tinghui Zhou, Shubham Tulsiani, Weilun Sun, Jitendra Malik, and Alexei A. Efros. View synthesis by appearance flow, 2016. 2, 8

[107] Xiangxin Zhu and Deva Ramanan. Face detection, pose estimation, and landmark localization in the wild. In *2012 IEEE conference on computer vision and pattern recognition*, pages 2879–2886. IEEE, 2012. 3

[108] Yi Zhu, Karan Sapra, Fitsum A. Reda, Kevin J. Shih, Shawn Newsam, Andrew Tao, and Bryan Catanzaro. Improving semantic segmentation via video propagation and label relaxation, 2018. 6