

Disentangling Propagation and Generation for Video Prediction

Hang Gao^{*1} Huazhe Xu^{*1} Qi-Zhi Cai^{†2} Ruth Wang^{†3} Fisher Yu¹ Trevor Darrell¹
 UC Berkeley¹ Sinovation Ventures AI Institute² Columbia University³

Abstract

A dynamic scene has two types of elements: those that move fluidly and can be predicted from previous frames, and those which are disoccluded (exposed) and cannot be extrapolated. Prior approaches to video prediction typically learn either to warp or to hallucinate future pixels, but not both. In this paper, we describe a computational model for high-fidelity video prediction which disentangles motion-specific propagation from motion-agnostic generation. We introduce a confidence-aware warping operator which gates the output of pixel predictions from a flow predictor for non-occluded regions and from a context encoder for occluded regions. Moreover, in contrast to prior works where confidence is jointly learned with flow and appearance using a single network, we compute confidence after a warping step, and employ a separate network to inpaint exposed regions. Empirical results on both synthetic and real datasets show that our disentangling approach provides better occlusion maps and produces both sharper and more realistic predictions compared to strong baselines.

1. Introduction

Video prediction is a challenge due to the many varied factors that combine to generate future appearance.

State-of-the-art approaches to video prediction are often purely *pixel-based* and generate each pixel from scratch (Figure 2a) [20, 38, 36, 37, 6, 43, 4]. They rely on 3D or recurrent convolutional networks to encode spatial contents over time, hoping to capture implicit motion representations. Another active line of purely *motion-based* approaches seeks pixel correspondence for explicit motion propagation [28, 19]; in this approach standard networks are augmented with an optical-flow representation and a warping function (Figure 2b). Because all pixels are copy-pasted from the history image buffer, a degree of temporal consistency and spatial richness are automatically ensured. However, disocclusions—locations in the target frame to which extrapolated optical flow has no projection—cause severe

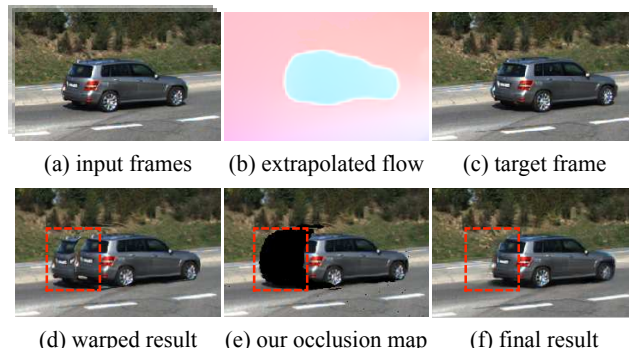


Figure 1: Video prediction by extrapolating the motion and hallucinating the pixels to be exposed. Given a video whose last frame shown in (a), our model first extrapolates optic flow (b) to its near future target (c). Direct motion propagation introduces ghosting effects (d) on pixels to be disoccluded (marked with the bounding box). Our disentangling approach computes a sparse occlusion map (e) and uses an contextual encoder to inpaint low-confidence patches, producing the final prediction (f).

errors with this family of models (Figure 1e).

We observe that humans can easily predict future appearance (e.g., in a scene such as shown in Figure 1) by tracking an object’s past motion and extrapolating future pixel values, as well as imagine the appearance of disoccluded pixels (the background) based on prior knowledge of similar scenes previously observed. We propose a compositional approach that disentangles spatial and temporal prediction, using a post-warp occlusion map to mediate between flow-based prediction and inpainting-based reconstruction of disoccluded regions.

Disentangling is different from the *linear fusion* approaches [9, 11] that use a single multi-task network to learn *pixel-* and *motion-based* predictions at the same time. To compose their results, the networks may also learn an extra occlusion map for weighted summation over each pixel (Figure 2c). Though those approaches are effective in many cases, it is unclear how much pixel and flow predictions can benefit each other. As our experiments indicate on the

^{*}Equal contribution.

[†]Work was done while Q.C. and R.W. were at Berkeley DeepDrive.

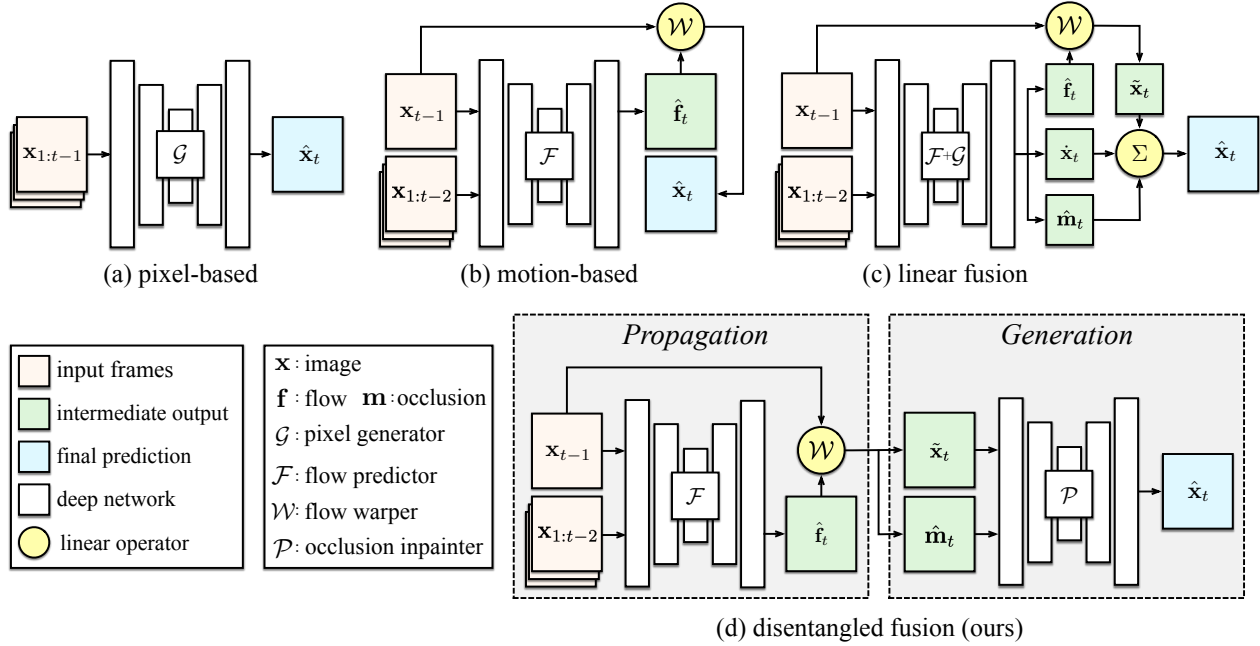


Figure 2: Overview of previous works and our proposed frameworks. Given $(t-1)$ frames and to predict one frame, previous video prediction methods typically consist of a flow predictor \mathcal{F} , or a pixel generator \mathcal{G} , or both: **(a)** *pixel-based* approaches; **(b)** *motion-based* approaches; **(c)** *linear-fusion* models; **(d)** ours.

real dataset, these two tasks do not actually correlate significantly with each other.

We propose a compositional framework that explicitly disentangles motion-specific propagation and motion-agnostic generation into two submodules, and thus called DPG, for high-quality video prediction. We factorize flow and pixel predictions into a serialized pipelines. We first employ a flow predictor for extrapolating optical flows and then an occlusion inpainter for context encoding based on images warped by flows. After flow prediction and warping, we compute confidence directly based on predicted flow rather than via joint unsupervised learning. The downstream occlusion inpainter generates pixels on occluded areas based on confidence map.

We evaluate our approach on both standard CalTech Pedestrian dataset [7] and more challenging KITTI Flow dataset [23] with larger motions and occlusions. Our approach achieves new state-of-the-art performance on both datasets with large performance gain on the perceptual realism metrics. To quantitatively evaluate our computed occlusion confidence maps, we perform ablation study on the RoamingImages [13] dataset where occlusion ground-truth masks are available. Our model again compare favorably against previous baselines in occlusion prediction in terms of Intersection over Union (IoU) metrics.

In summary, our main contributions are in two aspects: a

fusion pipeline that utilizes both optical flow and image synthesis for video prediction; a post-warp confidence-based occlusion map computation to disentangle pixel propagation and generation.

2. Related Works

2.1. Photo-realistic Image Synthesis

Realism is the constant pursuit of high-quality image synthesis [14, 5, 8, 42, 16, 44]. Recent developments towards photo-realistic image synthesis typically feature Generative Adversarial Networks (GANs) [10]. Conditioned on categorical labels [3], textual descriptions [31] or segmentations [40], high-fidelity images are shown to be able to synthesized. The closest work to ours are Dense Pose Transfer [25] and Transformation-grounded View Synthesis [26], which generally generate images by warping the original image with a learned appearance flow, conditioned on either given view transformation angle or a predefined dense pose, and then inpaint the ambiguous parts. In our case, however, the model needs to *predict* future motion and further synthesize based on both spatial and temporal information.

Spatial Context Encoding Pixels are not isolated. On the contrary, there are many cases [2, 1, 46] where appropriate spatial contexts must be retrieved. Spatial context encoders [27] query learned dataset priors with exposed appearance in search for missing patches. To our task, the

“mask” is an occlusion map where motion predictions are erroneous. Particularly, we employ partial convolutions [18] in our occlusion inpainter’s encoding blocks. Different from prior works on static images, our approach leverages not only the spatial but also temporal context.

2.2. High-fidelity Video Prediction

Different from image synthesis, video prediction not only cares for per-frame visual quality but also cross-frame consistency. Recent approaches [6, 4, 20, 29, 33] are typically *pixel-based*, which generate each pixel from scratch using implicit motion representations. Such methods may suffer from blurry effects, especially in the presence of unseen novel scenes. On the contrary, *motion-based* methods [28, 19] excel in making sharp predictions, yet fail in occlusion areas where motion predictions are erroneous or ill-defined. Meanwhile, Reda *et al.* [30] propose to model moving appearances with both convolutional kernels as in [9] and vectors as optical flow. Our closest prior work is [11] which also composes the pixel- and flow-based predictions through occlusion maps. However, our proposed method can differentiate in three aspects: (1) our pixel and flow prediction tasks are separately trained; (2) we employ an occlusion inpainter for our pixel generation so that more contextual information after warping can be utilized; (3) instead of predicting occlusion as another side task, we directly refer to predicted flows as the proxy for post warping confidence. It should be noted that [41] also seek for occlusion analytics to further parse flow reliability.

Disentangling Motion and Content Since videos are essentially moving contents whose semantics are agnostic to motions – one person’s identity does not change while he’s walking – it is natural to disentangle them apart. [36, 37, 6] are the three representative works in such direction. Though sharing the intuition to disentangle motion with other video property, our approach builds on lower-level vision such as pixel correspondence. We take advantage of optical flow and their embedded occlusion clues to naturally factorize pixel and flow predictions into separate modules.

3. Approach

Video prediction aims to synthesize future frames given a stack of history frames. For the ease of exposition, we here focus on a $(t - 1)$ -in 1-out prediction task: given an input video sequence denoted as $\mathbf{x}_{1:t-1}$, the model aims to predict the frame \mathbf{x}_t which should be accurate and visually sharp. In this paper, we build a model beyond simply using ℓ_1 or ℓ_2 reconstruction loss, which is widely known to encourage blurry results. Concretely, as illustrated in Figure 2d, our model is disentangled into two orthogonal yet complementary modules – a flow predictor \mathcal{F} and an occlusion inpainter \mathcal{P} to learn flow and pixel predictions separately.

Given the history frames $\mathbf{x}_{1:t-1}$, the flow predictor \mathcal{F} learns to predict the flow field $\hat{\mathbf{f}}_t$ for the pixel correspondence between the last input frame \mathbf{x}_{t-1} and the target \mathbf{x}_t . Utilizing $\hat{\mathbf{f}}_t$, we can propagating the input frame \mathbf{x}_{t-1} into a motion-dependent prediction $\tilde{\mathbf{x}}_t$. Since it neglects the existence of occlusions, we additionally compute an occlusion map $\hat{\mathbf{m}}_t$ during the propagation. Based on $\tilde{\mathbf{x}}_t$ and $\hat{\mathbf{m}}_t$, the inpainting module \mathcal{P} learns to inpaint occluded region from a learned prior within a training set. By explicitly disentangling future dynamics to motion-dependent areas and motion-agnostic areas, our model is able to predict high-fidelity future frames.

3.1. Motion Propagation

Our module \mathcal{F} computes the correlation of appearances between each pair of frames from the history inputs and predicts future motion dynamics using optical flow. We choose it over other motion representations, such as frame differences [12] or sparse trajectories [11] because it provides rich information about motion occlusions over pixels.

As illustrated in Figure 3a, the module \mathcal{F} is an encoder-decoder network with skip connections. The output of \mathcal{F} is a 2-channel flow field $\hat{\mathbf{f}}_t$ that aims to propagate the last frame \mathbf{x}_{t-1} into the predicted target frame $\tilde{\mathbf{x}}_t$. Formally, let $\{(i, j)\} \in \mathbf{x}_t$ be a Cartesian grid over the target frame, and we have

$$\hat{\mathbf{f}}_t = \{(\Delta_i, \Delta_j)\} = \mathcal{F}(\mathbf{x}_{1:t-1}). \quad (1)$$

By assuming local linearity, we can sample the future frame from the last given frame as

$$\tilde{\mathbf{x}}_t = \mathcal{S}(\mathbf{x}_{t-1}; \hat{\mathbf{f}}_t), \quad (2)$$

where $\mathcal{S}(\cdot; \hat{\mathbf{f}}_t)$ is a bilinear sampler that generates the new image by first mapping the regular grid to the transformed grid and then bilinearly interpolating between produced sub-pixels.

For training the module \mathcal{F} , we adopt a masked pixel loss \mathcal{L}_p and a smooth loss \mathcal{L}_{smt} which is similar to previous works on unsupervised flow estimation [22, 45]

$$\mathcal{L}_p(\tilde{\mathbf{x}}_t, \mathbf{x}_t; \hat{\mathbf{m}}_t) = \alpha \frac{1 - SSIM(\tilde{\mathbf{x}}_t \odot \hat{\mathbf{m}}_t, \mathbf{x}_t \odot \hat{\mathbf{m}}_t)}{2} + (1 - \alpha) \|\tilde{\mathbf{x}}_t \odot \hat{\mathbf{m}}_t - \mathbf{x}_t \odot \hat{\mathbf{m}}_t\|_1 \quad (3)$$

$$\mathcal{L}_{smt}(\hat{\mathbf{f}}_t, \mathbf{x}_t) = \sum_{i,j} |\nabla \hat{\mathbf{f}}_t(i, j)| \cdot (e^{-|\nabla \mathbf{x}_t(i, j)|})^T, \quad (4)$$

where $SSIM(\cdot, \cdot)$ denotes structural similarity index, \odot denotes element-wise product, ∇ is a vector differential operator, T denotes the transpose of image gradient weighting, and α is our trade-off weight between the loss terms, which is fixed at 0.9 through cross-validations.




Figure 3: The correlation between pixel and flow prediction tasks evaluated by different models. We show scatter plots about generation quality versus flow prediction error on KITTI Flow dataset [24]. Pixel and flow prediction tasks do not actually correlate significantly with each other on previous models using multi-task learning. Our disentangled fusion pipeline shows that after factorizing these tasks apart, both tasks could be better learned.

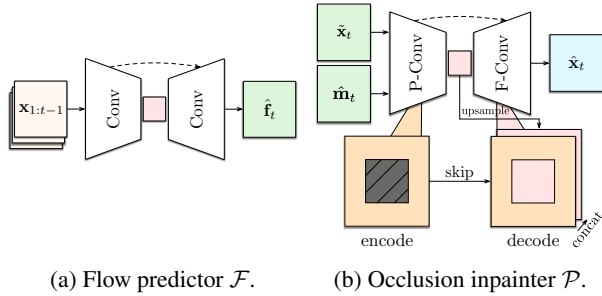


Figure 4: Conceptual illustrations of our \mathcal{F} and \mathcal{P} . (a) Our propagation module \mathcal{F} is a standard encoder-decoder fully convolutional network with skip connections, which takes in history frames $\mathbf{x}_{1:t-1}$ and predict the backward flow $\hat{\mathbf{f}}_t$ from \mathbf{x}_t to \mathbf{x}_{t-1} . (b) taking in the computed occlusion map $\hat{\mathbf{m}}_t$ and the warped image $\tilde{\mathbf{x}}_t$, our inpainting module \mathcal{P} inpaints for the final prediction $\hat{\mathbf{x}}_t$. It replaces standard convolutional blocks with partial convolutions (P-Conv) and fusion blocks (F-Conv) in its encoder and decoder, respectively.

The training loss for module \mathcal{F} is formulated as

$$\mathcal{L}_{\mathcal{F}} = \mathcal{L}_p(\tilde{\mathbf{x}}_t, \mathbf{x}_t; \hat{\mathbf{m}}_t) + \lambda_{smt} \mathcal{L}_{smt}(\hat{\mathbf{f}}_t, \mathbf{x}_t), \quad (5)$$

where λ_{smt} is a hyper-parameter that control the training schedule, we use $\lambda_{smt} = 0.1$ through a coarse grid search.

3.2. Occlusion Map

However, \mathcal{F} introduces “ghosting” effect in dis-occluded regions since regions in the target frame to which extrapolated optical flow has no projection (see Figure 1d and Figure 5 for detailed explanation). Intuitively, the “ghosting” area should be excluded from the propagated results by an

occlusion Map. It should be noted that similar ideas have been previously explored in [9, 11]; but in their contexts, occlusion maps are *linearly learned* to compose pixels from different sources. We argue that these regions can be *explicitly computed* based on backward flows, and hence can be directly masked out. We here describe the computation of occlusion map from an energy-based perspective. And we will later demonstrate the effectiveness of our computed map over the learned map in Section 4.4.

The pixel density can be viewed as an energy map: initially, it is uniformly filled for all pixels, and the motion propagation changes its distribution and make the energy map become dense or sparse on different region.

We initialize the energy field of the first frame to be a matrix filled with ones, denoted as $\mathbf{E}^1 = \mathbf{1}^{H \times W}$ for an image of $H \times W$ size.

Given a flow field, for each pixel in the first frame, the energy unit on each coordinate will be added into its 4 corresponded coordinates in the second frame bilinearly according to the flow field and we then get a new energy field \mathbf{E}^2 . We consider two special cases for each coordinate (x, y) in the second frame:

1. If $\mathbf{E}_{x,y}^2 = 0$, there is no pixel moving to this coordinate, which indicates it will be *dis-occluded*;
2. If $\mathbf{E}_{x,y}^2 > 2$, there are at least two pixels in the first frame compete for the same location, which suggests it will be *occluded*.

After complete definition of occluded and dis-occluded regions, we can then get the occlusion map according to

$$\hat{\mathbf{m}}_{i,j} = \begin{cases} 1 & 0 < \mathbf{E}_{i,j}^2 < 2, \\ 0 & \text{otherwise.} \end{cases} \quad (6)$$

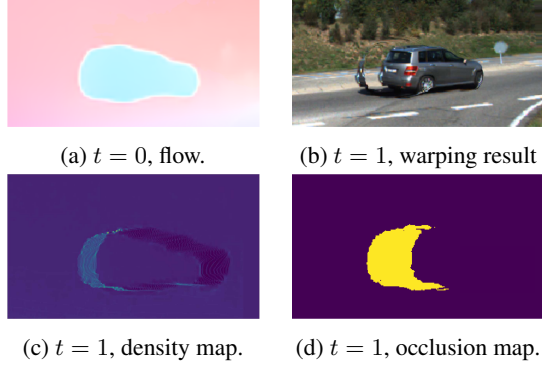


Figure 5: Ghosting effect caused by warping. Consider a foreground object on the background. (a) shows a predicted optical flow $\hat{\mathbf{f}}_1$ from target frame to current frame. (b) shows ghosting effects on the propagated frame $\tilde{\mathbf{x}}_1$ on locations to which flow has no projection. (c) shows our pixel density map computed by procedure described by Section 3.2. (d) shows our sparse occlusion map on which occlusion locations are colored as yellow.

3.3. Occlusion Inpainting as Context Encoding

Given the propagated frame $\tilde{\mathbf{x}}_t$ and the computed occlusion map $\hat{\mathbf{m}}_t$, we can now formulate our second modeling stage as context encoding to inpaint the missing pixels that are left out after propagation module. Our inpainting module \mathcal{P} adopts generally the same network architecture as its propagation counterpart with partial convolution blocks [18]. As illustrated in Figure 3b, our encoder takes the previous propagated frame $\tilde{\mathbf{x}}_t$ and its occlusion map $\hat{\mathbf{m}}_t$ as inputs, producing a latent feature representation. The decoder then takes this feature representation and synthesizes the missing content.

Specifically in the encoder, the partial convolution operators [18] mask out invalid pixels and re-normalize features within clean receptive fields only. For the extreme cases where all pixels in the receptive field are masked, we will simply return a zero value as the result. One important design choice is that the receptive field of the bottleneck should be bigger than the maximal area of the occlusion masks so that the feature map can be free from mask in our bottleneck.

For the decoder, the k th layer feature maps are upsampled and linked with the k th layer features from encoder counting from the latent layer in reverse order by skip connections. However, this raises a fusion issue that the feature from the encoder is from an image with invalid pixels. Previously in [18], feature maps and masks are concatenated channel-wisely and handled by new partial convolutions in their decoder. We find this could be improved by directly refilling the occluded encoder features by the upsampled decoder features. This decoding fusion is repeated at each

layer from the bottleneck up to our final output $\hat{\mathbf{x}}_t$.

To train our inpainting module \mathcal{P} , we design all of our losses to be temporal-independent so that the module can focus on the visual quality. In general, our loss terms consists of the following terms.

1. The pixel reconstruction loss

$$\mathcal{L}_{pix} = \mathcal{L}_p(\hat{\mathbf{x}}_t, \mathbf{x}_t; \hat{\mathbf{m}}_t) + \beta \mathcal{L}_p(\hat{\mathbf{x}}_t, \mathbf{x}_t; 1 - \hat{\mathbf{m}}_t),$$

where \mathcal{L}_p is defined in Equation 3.

2. The perceptual and style losses in VGG's n -dimensional latent space $\{\Psi^n\}$ as in [34].

$$\begin{aligned} \mathcal{L}_{prc} &= \frac{1}{n} \sum_n \|\psi(\tilde{\mathbf{x}}) - \psi(\mathbf{x})\| \odot \hat{\mathbf{m}}_t \|_1 \\ &\quad + \beta \|\psi(\tilde{\mathbf{x}}) - \psi(\mathbf{x})\| \odot (1 - \hat{\mathbf{m}}_t) \|_1, \\ \mathcal{L}_{sty} &= \frac{1}{n} \sum_n \|(\psi(\tilde{\mathbf{x}}) - \psi(\mathbf{x}))^T (\psi(\tilde{\mathbf{x}}) - \psi(\mathbf{x})) \odot \hat{\mathbf{m}}_t \|_1 \\ &\quad + \beta \|(\psi(\tilde{\mathbf{x}}) - \psi(\mathbf{x}))^T (\psi(\tilde{\mathbf{x}}) - \psi(\mathbf{x})) \odot (1 - \hat{\mathbf{m}}_t) \|_1. \end{aligned}$$

3. The total-variance loss to encourage similar texture in occlusion boundaries

$$\mathcal{L}_{var} = \sum_{i,j} \sqrt{\|\mathbf{x}_{t(i,j+1)} - \mathbf{x}_{t(i,j)}\|_2^2 + \|\mathbf{x}_{t(i+1,j)} - \mathbf{x}_{t(i,j)}\|_2^2}.$$

4. The extra semantic loss to enforce layout consistency between segmentation masks, distilled by a pretrained segmentation network [32] $\Phi : \mathcal{X} \rightarrow \mathcal{Y}$, since unconditional image inpainting tends to remove the foreground objects,

$$\begin{aligned} \mathcal{L}_{seg} &= CE(\Phi(\hat{\mathbf{x}}_t) \odot \hat{\mathbf{m}}_t, \mathbf{y}_t \odot \hat{\mathbf{m}}_t) \\ &\quad + \beta CE(\Phi(\tilde{\mathbf{x}}_t) \odot (1 - \hat{\mathbf{m}}_t), \mathbf{y}_t \odot (1 - \hat{\mathbf{m}}_t)). \end{aligned}$$

The total loss for the inpainting module is

$$\begin{aligned} \mathcal{L}_{\mathcal{P}} &= \mathcal{L}_{pix} + \lambda_{prc} \mathcal{L}_{prc} + \lambda_{sty} \mathcal{L}_{sty} \\ &\quad + \lambda_{var} \mathcal{L}_{var} + \lambda_{seg} \mathcal{L}_{seg}. \end{aligned} \quad (7)$$

In the above losses, β is our attentive weight for masked regions and all λ s are fixed hyper-parameter during training. We use a coarse grid search to set $\lambda_{prc} = 0.05$, $\lambda_{sty} = 120$, $\lambda_{var} = 0.1$, and $\lambda_{seg} = 5$ and we find best experimental performance (see Table 5) when $\beta = 10$.

3.4. Training

The overall training objective is formulated as

$$\min_{\mathcal{F}, \mathcal{P}} (\mathcal{L}_{\mathcal{F}} + \mathcal{L}_{\mathcal{P}}). \quad (8)$$

Because flow estimation and prediction are hard to learn and sensitive to data biases, we first train our motion propagation module \mathcal{F} and inpainting module \mathcal{P} separately. After gradients become stable, we connect the two components together and fine-tune the whole network in an end-to-end fashion.

4. Experiments

In this section, we will first introduce our benchmark datasets and baselines from previous works. Then we will verify our model’s performance on next-frame predictions and multi-frame predictions. To better understanding how we disentangle propagation and generation through an occlusion inpainting way, we also investigate the properties of occlusion maps.

4.1. Experiment Setup

Metrics We adopt the existing Peak Signal-to-Noise Ratio (PSNR) and Structural Similarity Index Measure (SSIM) [42] metrics to measure the pixel/patch-wise accuracy. However, it is well-known that these metrics disagree with the human perception [43, 35, 15, 39] because they tend to encourage blurriness over naturalness. Therefore we also measure the realism of results by Learned Perceptual Image Patch Similarity (LPIPS¹) proposed by Zhang *et al.* [47]. Higher PSNR/SSIM scores and lower LPIPS distances suggest better performance.

Dataset We evaluate our results on the CalTech Pedestrian dataset [7] and the KITTI Flow dataset [23]. The previous setup on the CalTech Pedestrian dataset is to first train a model on the training split of KITTI Raw [24] proposed by Lotter *et al.* [20] and then directly test it on the testing set of CalTech Pedestrian. Frames are center-cropped and down-sampled to 128×160 pixels. Every 11 consecutive frames are divided and sampled as a training clip in which the first 10 frames are fed into the model as the input, and the 11th frame is used as prediction target. As the results, the training, validation and testing sets consists of 3738, 14, and 1948 clips.

The KITTI Flow dataset is originally designed as a benchmark for optical-flow evaluation. For it is featured with higher resolution, larger motion and more occlusions, we should mention that it is more challenging compared to the raw KITTI dataset. It contains 3823 examples for training, 378 for validation and 4167 for testing. For all input images, we down-sampled and then center-cropped them into 320×640 pixels. We apply data augmentation techniques such as random cropping and random horizontal flipping for all the models. In addition, we sample video clips of 5 frames (4-in 1-out) from the dataset using a sliding window. This amounts to 3500 clips for training and 4000 clips for testing.

Baselines We consider representative baselines from three model families for video prediction: (1) *pixel-based* methods, including Beyond MSE [21], PredNet [20], SVP-LP [6], MCNet [37], MoCoGAN [36], and ContextVP [4]; (2) *motion-based* method, including DVF [19]; (3) *linear fusion* methods, including Dual Motion GAN [17] and Ctrl-

Method	Family	PSNR \uparrow	SSIM \uparrow	LPIPS $\downarrow (\times 10^{-2})$
Repeat	—	23.3	0.779	5.11
BeyondMSE [21]	<i>P</i>	—	0.881	—
PredNet [20]	<i>P</i>	27.6	0.905	7.47
ContextVP [4]	<i>P</i>	28.7	0.921	6.03
DVF [19]	<i>M</i>	26.2	0.897	5.57
Dual Motion GAN [17]	<i>F</i>	—	0.899	—
CtrlGen [11]	<i>F</i>	26.5	0.900	6.38
DPG (Ours)	<i>F</i>	28.2	0.923	5.04

Table 1: Next-Frame Prediction results on CalTech Pedestrian. All models are trained on KITTI Raw dataset. The best results under each metric are marked in bold.

Gen [11]. Note that we have not included [9] in our experiments since on its proposed RobotPush datasets, PredNet and SVP-LP are stronger baselines. For SVP-LP, we sampled 100 trajectories for each clip as in the original paper. The reported results are the mean performance over all the prediction results.

4.2. Next-Frame Prediction

We begin our evaluations on next-frame prediction. The goal of our task is simple and intuitive: Given the input frames, we need to output the next frame as accurate as possible.

Our result on CalTech Pedestrian dataset is shown on Table 1. We compare our model against previous state-of-the-art methods on this dataset. Our model achieves comparable PSNR and SSIM scores with ContextVP [4]. Meanwhile, our method can predict non-stretching textures in those occluded regions, which leads to smaller perceptual dissimilarity measured by LPIPS. As shown in Figure 6, our model is robust for both pixel propagation and novel scene inference.

Apart from empirical improvements, we find that, in terms of LPIPS metric, all the evaluated state-of-the-art methods do no better than the most naive baseline — repeating the last input frame as the prediction. This suggests that the CalTech Pedestrian dataset consists of small motions that are not obvious for human perception. This motivates us to work on KITTI Flow dataset, which is more challenging so that learners can benefit from more inductive biases and thus be more robust.

Table 2 shows our results on KITTI Flow dataset. As resolution increases, previous pixel-based methods (PredNet, SVP-LP) suffer from a steeper learning curve and more uncertainty in the visual space, resulting in the noticeable drop in their performance. Though achieving the better pixel/patch accuracy, they underperform the weakest repeating baseline in terms of the perceptual similarity. Our DPG achieves the best results in all metrics, especially LPIPS, showing around 30% improvement over the second best result from DVF.

¹We use LPIPS Version 0.1 available online.

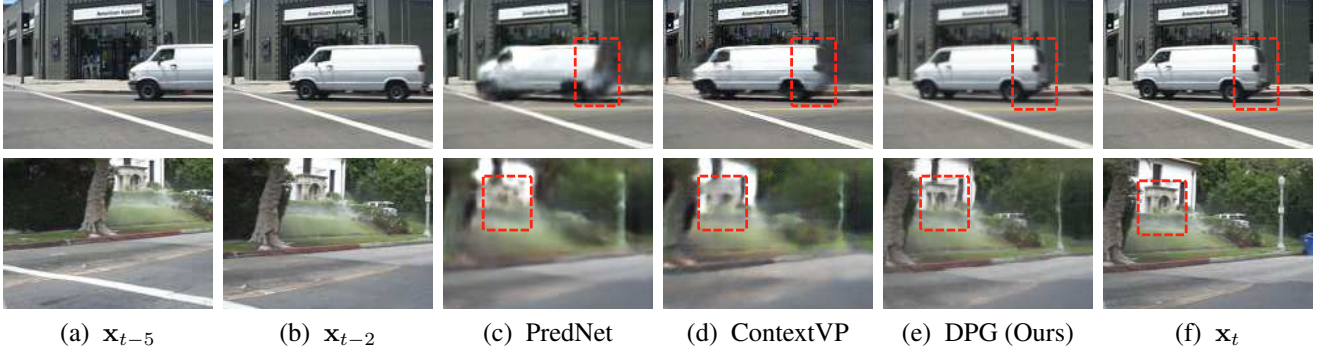


Figure 6: Qualitative comparisons for Next-Frame Prediction on CalTech Pedestrian dataset. Given 10 past frames, all models are required to predict the future frame in the next step. Our DPG produces much sharper results compared to previous state-of-the-art methods on this dataset. Remarkably, it is also capable of inferring semantically sensible appearance for occluded regions. Results are best viewed in color with zoom.

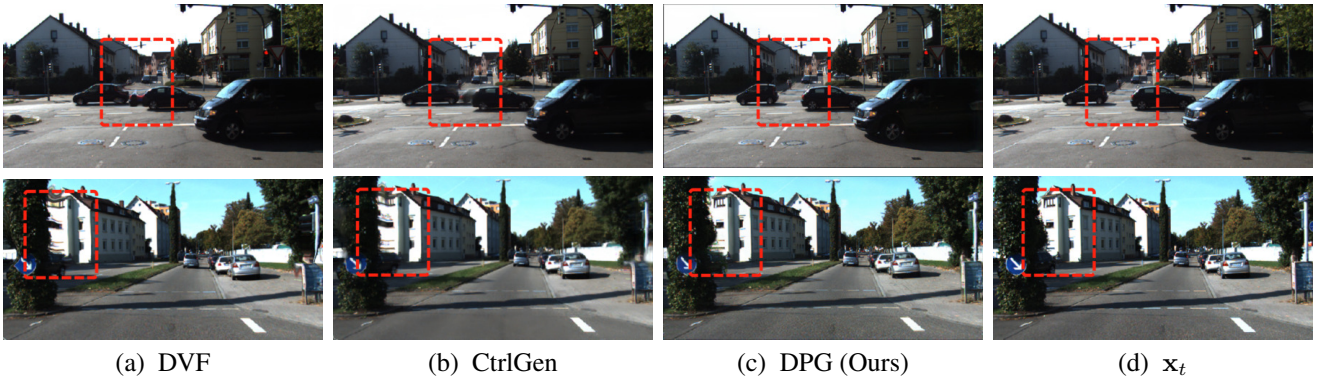


Figure 7: Qualitative comparisons for Next-Frame Prediction on the more challenging KITTI Flow dataset. All models are given 4 frames as input and required to predict the next frame. Frames include large motions, scene changes or occlusions. Our method is robust to these cases and consistent in the performance. More results please refer to our supplementary.

As demonstrated in Figure 7, our proposed model again produce more visually appealing predictions than our base-lines. In contrast to the pixel-based methods, all demonstrated methods suffer less from blurriness but display the distortion and stretch in shapes due to quick scene changes, which cause inaccurate flow prediction. Our model, instead, can predict better flow so as to alleviate undesirable artifacts in large motion areas. Occluded areas are masked by motion propagation and refilled by in-painting so that they are free of ghosting effects. Our occlusion in-painter learns a scene prior to hallucinate what is missing given the contextual information.

It is interesting to see that our model can achieve better PSNR/SSIM scores with training emphasized on realism terms (perceptual, style, and segmentation). We argue that this could serve as the evidence that our method could effectively learn dataset priors and flow prediction given the same amount of data.

Method	Family	PSNR \uparrow	SSIM \uparrow	LPIS $\downarrow (\times 10^{-2})$
Repeat	—	16.5	0.489	19.0
PredNet [20]	<i>P</i>	17.0	0.527	26.3
SVP-LP [6]	<i>P</i>	18.5	0.564	20.2
MCNet [37]	<i>P</i>	18.9	0.587	23.7
MoCoGAN [36]	<i>P</i>	19.2	0.572	18.6
DVF [19]	<i>M</i>	22.1	0.683	16.3
CtrlGen [11]	<i>F</i>	21.8	0.678	17.9
DPG (Ours)	<i>F</i>	22.3	0.696	11.4

Table 2: Next-Frame Prediction results on KITTI Flow. All models are trained to predict next frame given a history buffer of two frames. All evaluation results of the previous methods are obtained by their published codebases. The best results under each metric are marked in bold.

4.3. Multi-Frame Prediction

We next move to the more challenging Multi-frame Prediction task. Comparing to Next-Frame Prediction, it re-




Figure 8: Quantitative results for Multi-Frame Prediction on KITTI Flow dataset. All the models take in 4 frames as input and recursively predict next 8 frames. Best viewed in color with zoom. For qualitative results and SSIM quantitative result, please refer to our supplementary materials.

quires our models to output a sequence of frames to match the ground-truth future frames. To do a fair comparison between our method and other baselines, we limit all input frames to be 4 frames and requires each model to predict 8 following frames.

Figure 8 compares the results on Multi-Frame Prediction of our models with various baselines on KITTI Flow dataset. Our method shows consistent performance gains on all metrics through time. DVF performs similarly to our model for short-term prediction measured by PSNR but quickly decays after 2 steps. This is because that their method is sensitive to propagated error since there are no remedy mechanisms. Our model, however, can mask out undesirable regions and in-paint new pixels instead.

4.4. Ablations on Occlusion Map

To better evaluate the effectiveness of our computed occlusion map, we finally design a progressive curriculum for ablations.

Does occlusion map really help? We design different baselines to verify our design choices for computing the occlusion map. Results are shown on Table 3. The first baseline directly warps the last input frame as in [19], which also produces the result of our motion propagation module. Next baseline deploys an Auto-Encoder to directly refine warped frames without reasoning about occlusion map. The third one learns an independent occlusion map so as to fuse the result of motion propagation module and occlusion in-painter module, similar to [11], in an unsupervised way. It is interesting that the performance will downgrade if we only refine the motion propagated results without any guidance. Also, we find that without occlusion supervision in the wild, our computed occlusion map help final prediction performance better than its learned counterpart.

Why does occlusion map help? We are interested in how big is the gap between the their localization qualities. Since it is hard to get the ground truth occlusion map in the real world. We use a synthetic dataset instead. Here

Method	PSNR↑	SSIM↑	LPIPS↓
Warped only [19]	21.6	0.684	0.139
No occlusion map (auto-encoder)	21.3	0.674	0.120
Learned occlusion map [11]	21.8	0.679	0.179
Sparse occlusion map (Ours)	22.3	0.696	0.114

Table 3: Ablation study on different design choice for occlusion map.

Method	Learned occlusion map [11]	Sparse occlusion map (Ours)
IoU↑	0.0411	24.91

Table 4: Evaluation for learned occlusion map and our sparse occlusion map against ground-truth occlusion map.

Method	PSNR↑	SSIM↑	LPIPS↓
$\beta = 0$ (unmasked pixel only)	20.7	0.663	0.159
$\beta = 0.1$	20.9	0.671	0.134
$\beta = 1$ (same as unmasked pixel)	21.2	0.675	0.123
$\beta = \infty$ (masked pixel only)	18.4	0.514	0.206
$\beta = 10$ (Ours)	22.3	0.696	0.114

Table 5: Ablation study on different masking weight β .

we use RoamingImages [13] dataset for this specific quantitative study. The benchmark contains 80,000 examples as moving a random foreground image in front of a random background image according to a random linear motion. We report the IoU between the predicted occlusion map and the ground truth. Results are shown on Table 4. Not surprisingly, our computed occlusion map has much better quality compared to [11]. Another striking finding is that by also taking account of the results we shown in Table 3, the performance of learned occlusion maps is so bad that it can only marginally improve or even harm downstream prediction performance.

How does occlusion map help? We have shown the effectiveness of the computed occlusion map comparing to other design choices. Here we take a closer look to the importance of masking weight β in all of our training losses. As shown on Table 5, it is meaningful to use a larger loss weight for the occluded region during the training procedure.

5. Conclusion

In this work, we present a method for video prediction by disentangling motion-specific propagation and motion-agnostic generation. We propose a disentangled fusion pipeline which gates two tasks into two separate modules and directly evaluate pixel density map after warping for sparse occlusion maps. Experiments on synthetic and real datasets show our effectiveness against prior works.

References

- [1] Connelly Barnes, Eli Shechtman, Adam Finkelstein, and Dan B Goldman. Patchmatch: A randomized correspondence algorithm for structural image editing. *ACM Transactions on Graphics (ToG)*, 2009. 2
- [2] Marcelo Bertalmio, Guillermo Sapiro, Vincent Caselles, and Coloma Ballester. Image inpainting. In *Proceedings of the 27th annual conference on Computer graphics and interactive techniques*, 2000. 2
- [3] Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale gan training for high fidelity natural image synthesis. *arXiv preprint arXiv:1809.11096*, 2018. 2
- [4] Wonmin Byeon, Qin Wang, Rupesh Kumar Srivastava, and Petros Koumoutsakos. Contextvp: Fully context-aware video prediction. In *ECCV*, 2018. 1, 3, 6
- [5] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L Yuille. Semantic image segmentation with deep convolutional nets and fully connected crfs. *arXiv preprint arXiv:1412.7062*, 2014. 2
- [6] Emily Denton and Rob Fergus. Stochastic video generation with a learned prior. In *ECCV*, 2018. 1, 3, 6, 7
- [7] Piotr Dollar, Christian Wojek, Bernt Schiele, and Pietro Perona. Pedestrian detection: An evaluation of the state of the art. *IEEE transactions on pattern analysis and machine intelligence*, 2012. 2, 6
- [8] Alexey Dosovitskiy and Thomas Brox. Generating images with perceptual similarity metrics based on deep networks. In *NIPS*, 2016. 2
- [9] Chelsea Finn, Ian Goodfellow, and Sergey Levine. Unsupervised learning for physical interaction through video prediction. In *NIPS*, 2016. 1, 3, 4, 6
- [10] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *NIPS*, 2014. 2
- [11] Zekun Hao, Xun Huang, and Serge Belongie. Controllable video generation with sparse trajectories. In *CVPR*, 2018. 1, 3, 4, 6, 7, 8
- [12] De-An Huang, Vignesh Ramanathan, Dhruv Mahajan, Lorenzo Torresani, Manohar Paluri, Li Fei-Fei, and Juan Carlos Niebles. What makes a video a video: Analyzing temporal information in video understanding models and datasets. In *CVPR*, 2018. 3
- [13] Joel Janai, Fatma Güney, Anurag Ranjan, Michael J. Black, and Andreas Geiger. Unsupervised learning of multi-frame optical flow with occlusions. In *ECCV*, 2018. 2, 8
- [14] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *ECCV*, 2016. 2
- [15] Nal Kalchbrenner, Aaron van den Oord, Karen Simonyan, Ivo Danihelka, Oriol Vinyals, Alex Graves, and Koray Kavukcuoglu. Video pixel networks. *arXiv preprint arXiv:1610.00527*, 2016. 6
- [16] Chuan Li and Michael Wand. Combining markov random fields and convolutional neural networks for image synthesis. In *CVPR*, 2016. 2
- [17] Xiaodan Liang, Lisa Lee, Wei Dai, and Eric P Xing. Dual motion gan for future-flow embedded video prediction. In *ICCV*, 2017. 6
- [18] Guilin Liu, Fitsum A Reda, Kevin J Shih, Ting-Chun Wang, Andrew Tao, and Bryan Catanzaro. Image inpainting for irregular holes using partial convolutions. *arXiv preprint arXiv:1804.07723*, 2018. 3, 5
- [19] Ziwei Liu, Raymond A Yeh, Xiaoou Tang, Yiming Liu, and Aseem Agarwala. Video frame synthesis using deep voxel flow. In *ICCV*, 2017. 1, 3, 6, 7, 8
- [20] William Lotter, Gabriel Kreiman, and David Cox. Deep predictive coding networks for video prediction and unsupervised learning. In *ICLR*, 2017. 1, 3, 6, 7
- [21] Michael Mathieu, Camille Couprie, and Yann LeCun. Deep multi-scale video prediction beyond mean square error. *arXiv preprint arXiv:1511.05440*, 2015. 6
- [22] Simon Meister, Junhwa Hur, and Stefan Roth. Unflow: Unsupervised learning of optical flow with a bidirectional census loss. *arXiv preprint arXiv:1711.07837*, 2017. 3
- [23] Moritz Menze and Andreas Geiger. Object scene flow for autonomous vehicles. In *CVPR*, 2015. 2, 6
- [24] Moritz Menze, Christian Heipke, and Andreas Geiger. Joint 3d estimation of vehicles and scene flow. In *ISPRS Workshop on Image Sequence Analysis (ISA)*, 2015. 4, 6
- [25] Natalia Neverova, Rıza Alp Güler, and Iasonas Kokkinos. Dense pose transfer. *arXiv preprint arXiv:1809.01995*, 2018. 2
- [26] Eunbyung Park, Jimei Yang, Ersin Yumer, Duygu Ceylan, and Alexander C Berg. Transformation-grounded image generation network for novel 3d view synthesis. In *CVPR*, 2017. 2
- [27] Deepak Pathak, Philipp Krahenbuhl, Jeff Donahue, Trevor Darrell, and Alexei A Efros. Context encoders: Feature learning by inpainting. In *CVPR*, 2016. 2
- [28] Silvia L Pintea, Jan C van Gemert, and Arnold WM Smeulders. Déjà vu. In *ECCV*, 2014. 1, 3
- [29] MarcAurelio Ranzato, Arthur Szlam, Joan Bruna, Michael Mathieu, Ronan Collobert, and Sumit Chopra. Video (language) modeling: a baseline for generative models of natural videos. *arXiv preprint arXiv:1412.6604*, 2014. 3
- [30] 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. *arXiv preprint arXiv:1811.00684*, 2018. 3
- [31] Scott Reed, Zeynep Akata, Xinchun Yan, Lajanugen Logeswaran, Bernt Schiele, and Honglak Lee. Generative adversarial text to image synthesis. *arXiv preprint arXiv:1605.05396*, 2016. 2
- [32] Samuel Rota Bulò, Lorenzo Porzi, and Peter Kotschieder. In-place activated batchnorm for memory-optimized training of dnns. In *CVPR*, 2018. 5
- [33] Masaki Saito, Eiichi Matsumoto, and Shunta Saito. Temporal generative adversarial nets with singular value clipping. In *ICCV*, 2017. 3
- [34] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. 5

- [35] Nitish Srivastava, Elman Mansimov, and Ruslan Salakhudinov. Unsupervised learning of video representations using lstms. In *International conference on machine learning*, 2015. 6
- [36] Sergey Tulyakov, Ming-Yu Liu, Xiaodong Yang, and Jan Kautz. Mocogan: Decomposing motion and content for video generation. In *ICLR*, 2018. 1, 3, 6, 7
- [37] Ruben Villegas, Jimei Yang, Seunghoon Hong, Xunyu Lin, and Honglak Lee. Decomposing motion and content for natural video sequence prediction. In *ICLR*, 2017. 1, 3, 6, 7
- [38] Ruben Villegas, Jimei Yang, Yuliang Zou, Sungryull Sohn, Xunyu Lin, and Honglak Lee. Learning to generate long-term future via hierarchical prediction. In *ICML*, 2017. 1
- [39] Jacob Walker, Kenneth Marino, Abhinav Gupta, and Martial Hebert. The pose knows: Video forecasting by generating pose futures. In *ICCV*, 2017. 6
- [40] 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. In *CVPR*, 2018. 2
- [41] 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
- [42] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 2004. 2, 6
- [43] Nevan Wichers, Ruben Villegas, Dumitru Erhan, and Honglak Lee. Hierarchical long-term video prediction without supervision. In *ICML*, 2018. 1, 6
- [44] Saining Xie, Xun Huang, and Zhuowen Tu. Top-down learning for structured labeling with convolutional pseudoprior. In *ECCV*, 2016. 2
- [45] Zhichao Yin and Jianping Shi. Geonet: Unsupervised learning of dense depth, optical flow and camera pose. In *CVPR*, 2018. 3
- [46] Hang Zhang, Kristin Dana, Jianping Shi, Zhongyue Zhang, Xiaoang Wang, Amrith Tyagi, and Amit Agrawal. Context encoding for semantic segmentation. In *CVPR*, 2018. 2
- [47] 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. 6