

# Devon: Deformable Volume Network for Learning Optical Flow

Yao Lu  
Data61, ANU, ACRV

Jack Valmadre  
Google

Heng Wang  
Facebook

Juho Kannala  
Aalto University

Mehrtash Harandi  
Monash University

Philip H. S. Torr  
University of Oxford

## Abstract

*State-of-the-art neural network models estimate large displacement optical flow in multi-resolution and use warping to propagate the estimation between two resolutions. Despite their impressive results, it is known that there are two problems with the approach. First, the multi-resolution estimation of optical flow fails in situations where small objects move fast. Second, warping creates artifacts when occlusion or dis-occlusion happens. In this paper, we propose a new neural network module, Deformable Cost Volume, which alleviates the two problems. Based on this module, we designed the Deformable Volume Network (Devon) which can estimate multi-scale optical flow in a single high resolution. Experiments show Devon is more suitable in handling small objects moving fast and achieves comparable results to the state-of-the-art methods in public benchmarks.*

## 1. Introduction

Optical flow estimation is the problem of finding dense correspondences between two images. It is a classic problem in computer vision and has been studied for more than 30 years. Since Horn and Schunck’s variational method [14], a large number of optical flow algorithms [1, 8, 9] has been developed and lead to various applications such as tracking [11] and video classification [41]. Nevertheless, the problem of estimating optical flow is not yet solved. One can claim that even today we still do not have a practical algorithm which is fast, robust and accurate for real-world images.

Recently, supervised learning of optical flow with neural networks has shown great promises [7, 18, 31, 38, 39, 37, 15]. By making use of graphics engines, large scale datasets with synthetic images and ground-truth optical flow can be generated [4, 7, 28], which in return enable us to learn optical flow in an end-to-end manner. We note that while the training images are synthetic, neural networks seem to gen-

eralize reasonably well on real-world images [27]. Compared to classical methods, the neural network approaches have the advantage of offering robust features and fast inference [7, 18, 37, 15].

One of the major difficulties in estimating optical flow is large displacements. Learning optical flow with large displacements may not be achieved by Convolutional Neural Networks (CNNs) with only standard convolution operations. This is due to the fact that CNNs make use of small filters. It is prohibitive to increase filter sizes naïvely to cover large displacements, as the number of parameters and the computational cost both increase drastically.

To handle large displacements, multi-resolution estimation of optical flow is employed in several neural network models. SpyNet [31] downsamples the original images into multiple resolution levels and each level is handled by a CNN decoder to output optical flow of the corresponding resolution. PWC-Net [37] and LiteFlowNet [15] follow the same strategy except that they use strided convolutions instead of downsampling to reduce the resolution and use cost volumes as a more explicit representation of motion. As a result, a decoder in lower resolution has effectively a larger receptive size to cover large displacements.

However, as pointed out in [3, 36], the multi-resolution estimation of optical flow faces the “small objects move fast” problem as small objects disappear in lower resolution and cannot be recovered in higher resolution due to their large motion. Although the problem is not severe in terms of performance on the current public benchmarks, it limits the use of optical flow for high precision applications where small objects might contain vital information. The key to solve the problem is to handle multi-scale motion in a single high resolution [44]. Another technique to handle large displacements in neural network models such as [18, 31, 37] is warping, which propagates optical flow between two stages in a model. However, warping often creates distortions and artifacts. This issue will be discussed in details in §2.

In this paper, we introduce the Deformable Volume Network (Devon), which avoids the drawbacks of multi-resolution estimation and warping.

## 2. The Problem of Warping

Warping has been used in variational methods [3, 26] and neural network models [18, 31, 37] for iteratively refining optical flow estimations in a multi-stage framework. The first stage covers large displacements and outputs a rough estimation. Then the second image (or its feature maps) is warped by the roughly estimated optical flow such that pixels of large displacements in the second image are moved closer to their correspondences in the first image. As a result, the next stage, which receives the original first image and the warped second image as inputs, only needs to handle smaller displacements to refine the estimation.

Let  $I : \mathbb{R}^2 \rightarrow \mathbb{R}^3$  denote the first image,  $J : \mathbb{R}^2 \rightarrow \mathbb{R}^3$  denote the second image and  $F : \mathbb{R}^2 \rightarrow \mathbb{R}^2$  denote the optical flow field of the first image. The warped second image is defined as

$$\tilde{J}(\mathbf{p}) = J(\mathbf{p} + F(\mathbf{p})) \quad (1)$$

for image location  $\mathbf{p} \in \mathbb{R}^2$  [18].

The warping operation creates a transformed image reasonably well if the new pixel locations  $\mathbf{p} + F(\mathbf{p})$  do not occlude or collide with each other. For example, this is the case with the affine transform  $F(\mathbf{p}) = \mathbf{A}\mathbf{p} + \mathbf{t}$  where  $\mathbf{A}$  and  $\mathbf{t}$  are the transformation parameters. However, for real-world images, occlusions are common (*e.g.*, when an object moves and the background is still). If an image is warped with the optical flow which induces occlusions, duplicates will be created. The effect is demonstrated in Figure 1. The artifacts cannot be cleaned simply by subtracting the first or the second image from the warped image, as shown in Figure 1 (e) and (f). The artifacts induced by warping have been previously observed in [2, 42, 21]. Intuitively, if a pixel which is moved by warping to a new location and no other pixel are moved to fill in its old location, the pixel will appear twice in the warped image.

Mathematically, consider the following example. Assume the value of  $J(\mathbf{p}_1)$  is unique in  $J$ , that is,  $J(\mathbf{p}) \neq J(\mathbf{p}_1)$  for all  $\mathbf{p} \neq \mathbf{p}_1$ . Then for an optical flow field in which

$$F(\mathbf{p}_1) = 0, \quad F(\mathbf{p}_2) = \mathbf{p}_1 - \mathbf{p}_2, \quad (2)$$

we have

$$\tilde{J}(\mathbf{p}_1) = J(\mathbf{p}_1 + F(\mathbf{p}_1)) \quad (3)$$

$$= J(\mathbf{p}_1 + 0) = J(\mathbf{p}_1), \quad (4)$$

$$\tilde{J}(\mathbf{p}_2) = J(\mathbf{p}_2 + F(\mathbf{p}_2)) \quad (5)$$

$$= J(\mathbf{p}_2 + \mathbf{p}_1 - \mathbf{p}_2) = J(\mathbf{p}_1). \quad (6)$$

Therefore  $\tilde{J}(\mathbf{p}_1) = \tilde{J}(\mathbf{p}_2) = J(\mathbf{p}_1)$ . Since the value of  $J(\mathbf{p}_1)$  is unique in image  $J$  but not unique in  $\tilde{J}$ , a duplicate is created on the warped second image  $\tilde{J}$ .

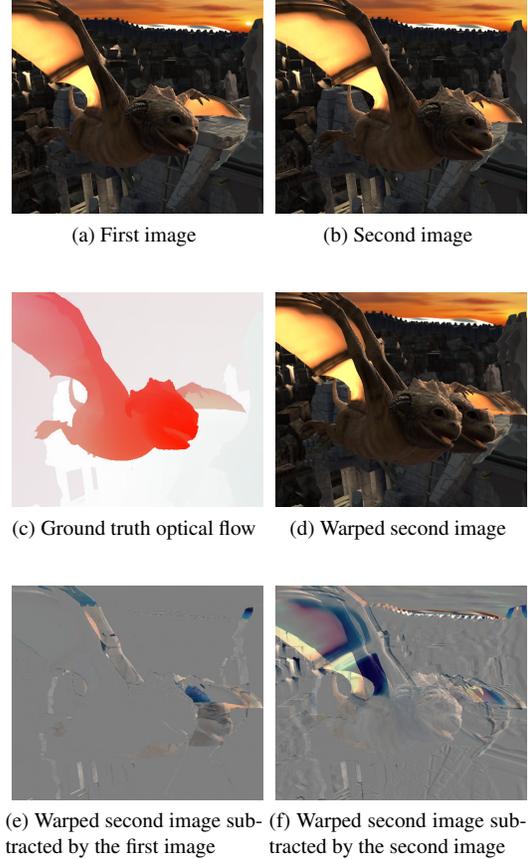


Figure 1: Artifacts of using image warping. From (d), we can see the duplicates of the dragon head and wings. The images and the ground truth optical flow are from the Sintel dataset [4]. Warping is done with function `image.warp()` in the Torch-image toolbox.

When the duplicates happen, it makes the optical flow estimation erroneous since artificial candidate correspondences are created. Although neural networks as universal approximators might be able to learn the self-corrected correspondences if trained with ground-truth optical flow, one should be aware of the issue which might cause problems in designing non-learning-based methods or more interpretable neural network models.

## 3. Deformable Cost Volume

Let  $I$  denote the first image,  $J$  denote the second image and  $f_I : \mathbb{R}^2 \rightarrow \mathbb{R}^d$  and  $f_J : \mathbb{R}^2 \rightarrow \mathbb{R}^d$  denote their feature maps of dimensionality  $d$ , respectively. The standard cost volume is defined as

$$C(\mathbf{p}, \mathbf{v}) = \|f_I(\mathbf{p}) - f_J(\mathbf{p} + \mathbf{v})\|, \quad (7)$$

for image location  $\mathbf{p} \in \mathbb{R}^2$ , neighbor  $\mathbf{v} \in [-\frac{k-1}{2}, \frac{k-1}{2}]^2$  of neighborhood size  $k$  and a given vector norm  $\|\cdot\|$ .

The cost volume gives an explicit representation of displacements. The idea of using cost volume goes back to stereo matching [35]. When using the feature maps learned by neural networks, construction and processing of a fully connected cost volume, in which the neighborhood is large enough to cover the maximum displacement, leads to high performance in stereo matching [46] and optical flow [43]. However, the displacements in stereo matching are one-dimensional while in optical flow they are two-dimensional. For two images (and their feature maps) of resolution  $m \times n$ , the construction of the cost volume in equation (7) has time and space complexity  $O(mndk^2)$ . Naïvely increasing neighborhood size  $k$  to cover large displacements increases the computation quadratically. As a result, DCFlow requires several seconds to compute optical flow for a pair of images on a GPU and large memory usage [43].

To reduce the computational burden, in [37, 15], multi-resolution feature maps of two images are created and a cost volume of a small neighborhood is constructed at each resolution. Warping is used to propagate the optical flow between two resolutions. However, as discussed before, the multi-resolution estimation of optical flow leads to the “small objects move fast” problem and warping induces artifacts and distortion. To avoid the drawbacks, we propose a new neural network module, the deformable cost volume. The key idea is: instead of deforming images or their feature maps, we deform the cost volume and leave the images and the feature maps unchanged.

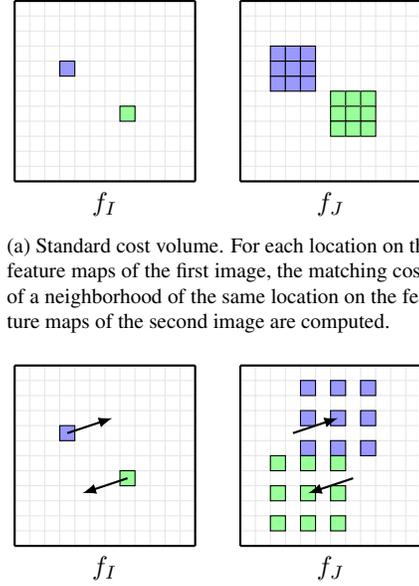
The proposed deformable cost volume is defined as

$$C(\mathbf{p}, \mathbf{v}, r, F) = \|f_I(\mathbf{p}) - f_J(\mathbf{p} + r \cdot \mathbf{v} + F(\mathbf{p}))\| \quad (8)$$

where  $r$  is the dilation rate and  $F(\cdot)$  is an external flow field. The dilation rate  $r$  is introduced to enlarge the size of the neighborhood to handle large displacements without reducing the resolution. This is inspired by the dilated convolution [5, 45] which enlarges its receptive field in a similar way.  $F(\cdot)$  can be obtained from the optical flow estimated from a previous stage or an external algorithm. If  $F(\mathbf{p}) = \mathbf{0}$  for all  $\mathbf{p}$  and  $r = 1$ , then the deformable cost volume is reduced to the standard cost volume. For non-integer  $F(\mathbf{p})$ , bilinear interpolation is used. The deformable cost volume is illustrated in Figure 2.

Since the deformable cost volume does not distort  $f_I$  or  $f_J$ , the artifacts associated with warping will not be created. Optical flow can be inferred from the deformable cost volume solely without resorting to the feature maps of the first image to counter the duplicates.

The deformable cost volume is differentiable with respect to  $f_I(\mathbf{p})$  and  $f_J(\mathbf{p} + r \cdot \mathbf{v} + F(\mathbf{p}))$  for each image location  $\mathbf{p}$ . Due to bilinear interpolation, the deformable cost volume is also differentiable with respect to  $F(\mathbf{p})$ , using the same technique as in [18, 20]. Therefore, the deformable cost volume can be inserted in a neural network



(a) Standard cost volume. For each location on the feature maps of the first image, the matching costs of a neighborhood of the same location on the feature maps of the second image are computed.

(b) Deformable cost volume. For each location on the feature maps of the first image, the matching costs of a **dilated** neighborhood of the same location, **offset by a flow vector**, on the feature maps of the second image are computed.

Figure 2: Cost Volumes

for end-to-end learning optical flow.

To see how the deformable cost volume avoids the artifacts of warping, consider the following. Assume  $F(\mathbf{p}_1) = \mathbf{0}$  and  $F(\mathbf{p}_2) = \mathbf{p}_1 - \mathbf{p}_2$ . The standard cost volume (7) with warping (1) gives

$$C(\mathbf{p}_1, \mathbf{0}) = \|f_I(\mathbf{p}_1) - \tilde{f}_J(\mathbf{p}_1 + \mathbf{0})\| \quad (9)$$

$$= \|f_I(\mathbf{p}_1) - f_J(\mathbf{p}_1 + F(\mathbf{p}_1))\| \quad (10)$$

$$= \|f_I(\mathbf{p}_1) - f_J(\mathbf{p}_1)\| \quad (11)$$

$$C(\mathbf{p}_1, \mathbf{p}_2 - \mathbf{p}_1) = \|f_I(\mathbf{p}_1) - \tilde{f}_J(\mathbf{p}_1 + \mathbf{p}_2 - \mathbf{p}_1)\| \quad (12)$$

$$= \|f_I(\mathbf{p}_1) - \tilde{f}_J(\mathbf{p}_2)\| \quad (13)$$

$$= \|f_I(\mathbf{p}_1) - f_J(\mathbf{p}_2 + F(\mathbf{p}_2))\| \quad (14)$$

$$= \|f_I(\mathbf{p}_1) - f_J(\mathbf{p}_1)\| \quad (15)$$

That  $C(\mathbf{p}_1, \mathbf{0}) = C(\mathbf{p}_1, \mathbf{p}_2 - \mathbf{p}_1)$  implies  $f_I(\mathbf{p}_1)$  has the same matching cost for  $f_J(\mathbf{p}_1)$  and  $f_J(\mathbf{p}_2)$ , which does not hold in general and makes the matching ambiguous. On the other hand, with deformable cost volume (8) of dilation rate one, we have

$$C(\mathbf{p}_1, \mathbf{0}) = \|f_I(\mathbf{p}_1) - f_J(\mathbf{p}_1)\| \quad (16)$$

$$C(\mathbf{p}_1, \mathbf{p}_2 - \mathbf{p}_1) = \|f_I(\mathbf{p}_1) - f_J(\mathbf{p}_2 + F(\mathbf{p}_1))\| \quad (17)$$

$$= \|f_I(\mathbf{p}_1) - f_J(\mathbf{p}_2)\| \quad (18)$$

As  $C(\mathbf{p}_1, \mathbf{0}) \neq C(\mathbf{p}_1, \mathbf{p}_2 - \mathbf{p}_1)$  in general, the artifact is avoided.

## 4. Deformable Volume Network

Our proposed model is the Deformable Volume Network (Devon), as illustrated in Figure 3. Devon has multiple stages. Each stage is a neural network with an identical Siamese architecture, which consists of an encoding module, a relation module and a decoding module. Each stage outputs the optical flow field of 1/4 resolution and then the flow from last stage is bilinearly upsampled to obtain the final prediction. The optical flow estimated from a previous stage is propagated to the current one through the deformable cost volume and residual connections.

Compared to previous neural network models [7, 18, 31, 37, 15], Devon is different in the following major ways: (1) All stages in Devon output optical flow of the same resolution. While Devon handles multi-scale motion by the deformable cost volume, it does not use the multi-resolution representation of motion. Extensive downsampling leads to the loss of information and deficiency in handling the “small objects move fast” problem. The advantage of estimating motion direct on a single high resolution is also shown in [44]. (2) Each stage acts on the undistorted images. No warping is used. Therefore, the artifacts discussed in §2 can be avoided. (3) The decoding module only receives inputs from the relation module. Therefore, neural networks infer the optical flow solely from the relations between two images, rather than memorize the optical flow pattern of a single image as a short-cut. The short-cut issue has appeared when applying neural networks to learn monocular stereo [40]. On the contrary, in FlowNetC, PWC-Net and LiteFlowNet, the decoding module also receives inputs from the encoding module of the first image. (4) The encoding module is shared in all stages.

We describe the details of each module structure below. All convolution layers have zero-padding size one. Besides, all convolution layers, except the last one in the encoding module and the last one in the decoding module, are followed by a leaky ReLU function [12] with leakiness 0.1.

### 4.1. Encoding Module

The encoding module has a U-Net structure [33] with residual connections [13], as shown in Figure 4. Since the model output optical flow of 1/4 resolution, the encoding module has 6 convolution layers of stride 2 but 4 upsampling layers. We also experimented with a simpler module and result is shown in the ablation analysis in supplementary materials.

### 4.2. Relation Module

The relation module is illustrated in Figure 5. It concatenates the outputs of five deformable cost volumes, which have different neighborhood size  $k$  or dilation rate  $r$ , as shown in Table 1. Such combination enables dense correspondences nearby the center of an image location and

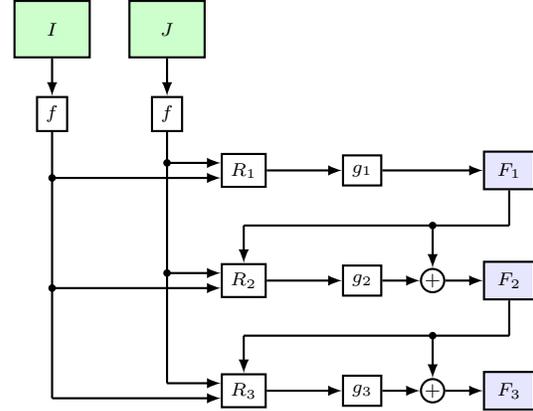


Figure 3: Deformable Volume Network (Devon) with three stages.  $I$  denotes the first image,  $J$  denotes the second image,  $f$  denotes the encoding module (§4.1),  $R_t$  denotes the relation module (§4.2),  $g_t$  denotes the decoding module (§4.3) and  $F_t$  denotes the estimated optical flow for stage  $t$ .

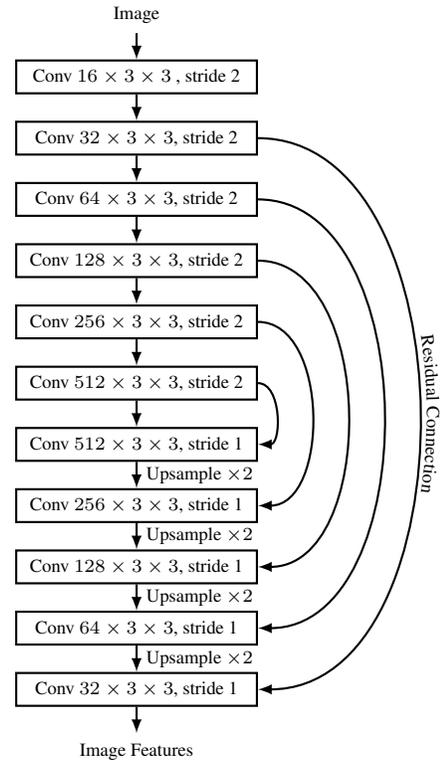


Figure 4: Encoding module  $f$ . The residual connection denotes the output of a layer is added to the output of another layer.

sparse correspondences in peripheral to capture multi-scale motion. This is consistent with the fact that small displacements are more frequent in natural videos [34] and resem-

bles the structure of retina, as illustrated in Figure 6.

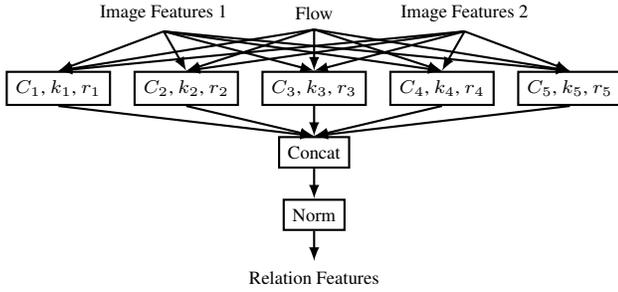


Figure 5: Relation module  $R$ .  $C_1 \sim C_5$  denote the deformable cost volumes.  $k_1 \sim k_5$  denote the neighborhood sizes.  $r_1 \sim r_5$  denote the dilation rates. Concat denotes concatenation. Norm denotes normalization.

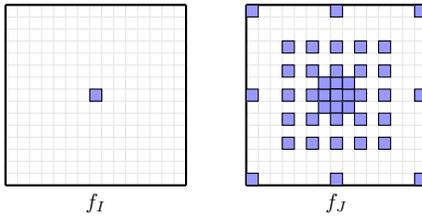


Figure 6: Concatenation of deformable cost volumes creates a retinal structure of correspondences. In this example, three cost volumes of neighborhood sizes  $(k_1, k_2, k_3) = (3, 5, 3)$  and dilation rates  $(r_1, r_2, r_3) = (1, 2, 7)$  respectively are concatenated.

	$(k_1, k_2, k_3, k_4, k_5)$	$(r_1, r_2, r_3, r_4, r_5)$
$R_1$	(5, 5, 5, 5, 9)	(1, 3, 8, 12, 20)
$R_2$	(5, 5, 5, 5, 9)	(1, 3, 8, 10, 12)
$R_3$	(5, 5, 5, 5, 9)	(1, 3, 4, 5, 7)

Table 1: Hyperparameters of deformable cost volumes in Devon.

Since Devon is a multi-stage model which performs coarse-to-fine estimation of optical flow, we gradually decrease the dilation rates as the optical flow is expected to get finer in later stages. We also experimented with using the cost volumes of the same neighbor sizes and dilation rates for all three stages, the result of which is reported in supplementary materials.

Next, for each location in the concatenated feature maps, a normalization method is applied across the channels. We apply  $f(C) = \exp(-C)$  elements-wise, which encourages the elements representing the minimum cost to pop-up. Such normalization improves the estimation accuracy

as shown in §5.3.

The output of this module has size  $(k_1^2 + k_2^2 + k_3^2 + k_4^2 + k_5^2) \times m \times n$ , where  $m$  is the height and  $n$  is the width of the feature maps. When the module does not receive an optical flow as one of the inputs (in the first stage), it is set to receive a zero-valued optical flow field.

Since the relation module captures multi-scale motion without reducing the resolution, the presence and precise location of small objects which move fast are retained. This allows Devon to have better chance in solving the “small objects move fast” problem.

### 4.3. Decoding Module

In the decoding module, we again use the U-Net structure with residual connections. The whole module structure is illustrated in Figure 7. Each stage has its own decoder. We also experimented with sharing decoder in all three stages. The result is reported in supplementary materials.

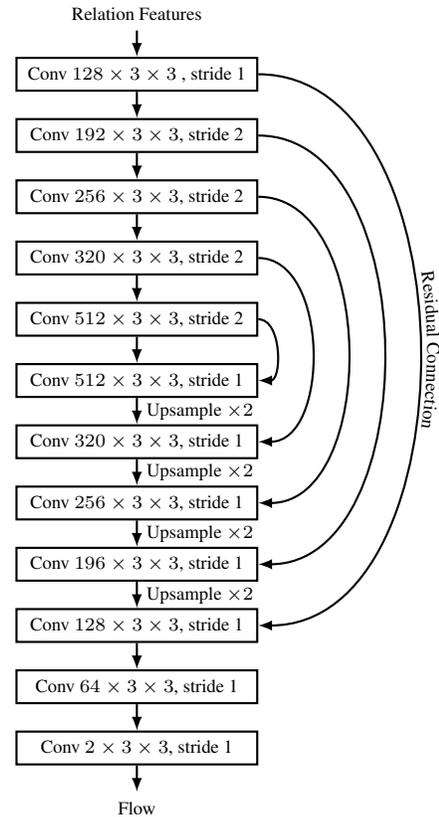


Figure 7: Decoding module  $g$ . The residual connection denotes the output of a layer is added to the output of another layer.

## 5. Experiments

We evaluate Devon on two challenging benchmarks which contain large motions: Sintel [4] and KITTI [10], as in [7, 31, 18, 37]. We compare Devon with the state-of-the-art neural network models: PWC-Net [37] and LiteFlowNet [15]. We use Devon with three stages and  $l_1$  norm for the deformable cost volumes.

### 5.1. Training

Our training procedure largely follows from [18, 37]. We first train our network on FlyingChairs. We use the  $l_2$  loss function

$$L = \sum_t \gamma_t |F_{GT} - \hat{F}_t|_2 \quad (19)$$

where  $F_{GT}$  denotes the ground-truth optical flow and  $\hat{F}_t$  denotes the bilinear upsampled network output at stage  $t$ . For Devon with three stages, we choose  $\gamma_1 = 0.2$ ,  $\gamma_2 = 0.3$  and  $\gamma_3 = 0.5$ . All pixel values of the images are multiplied by  $1/255$ . Empirically, such normalization is found to accelerate the training. For optimization algorithm, we use Adam [23] with  $\beta_1 = 0.9$ ,  $\beta_2 = 0.999$  and weight decay factor 0.0004. We use the weight initialization method in [12]. We use data augmentation which is consisted of random cropping of size  $448 \times 384$ , translation, rotation, color jittering and Gaussian noise. There are totally 500000 mini-batch updates with mini-batch size 8. After training on FlyingChairs, we fine-tune our model on FlyingThings3D.

For the evaluation on Sintel, we fine-tuned the model, which was previously trained on FlyingChairs and then fine-tuned on FlyingThings3D, on the Sintel training set. For the evaluation on KITTI, we fine-tuned the model, which was previously trained on FlyingChairs and then fine-tuned on FlyingThings3D, on the KITTI training set.

All the experiments are conducted with PyTorch. The deformable cost volume is implemented in CUDA with PyTorch interface.

### 5.2. Main Results

	Sintel Clean	Sintel Final	KITTI 2012	KITTI 2015
	EPE	EPE	EPE	F1-all
PWC-Net	4.39	<b>5.04</b>	1.7	<b>9.16 %</b>
LiteFlowNet	4.54	5.38	<b>1.6</b>	9.38%
Devon	<b>4.34</b>	6.35	2.6	14.31 %

Table 2: Results on the whole Sintel and KITTI datasets.

In Figure 8, 9 and 10, we show visualization results of situations where small objects move fast. All models were trained on FlyingChairs and then fine-tuned on FlyingThings3D. No additional fine-tuning is applied. From Figure 8, 9 and 10, we can see Devon gives much more accurate estimation of the small objects.

The results on whole Sintel and KITTI are in Table 2. Note that our goal is not to outperform previous methods on the whole Sintel and KITTI datasets, which contain mostly large objects. We aim to solve a long-lasting problem in optical flow estimation, “small objects moving fast”. For future work, our model can be integrated with multi-resolution models to handle objects with various sizes and speed.

### 5.3. Ablation Analysis

We perform an ablation analysis of Devon by replacing the deformable cost volume with warping and standard cost volume. The models were trained on FlyingChairs. The timing was recorded on a NVIDIA TITAN Xp graphics card for processing a pair of RGB images of size  $1024 \times 448$ . The results are listed in Table 3 and 4. From the results, we can see deformable cost volume leads to higher speed and comparable accuracy. This is because the deformation and the cost volume construction are merged into one process in deformable cost volume and therefore reduces the runtime. Therefore, deformable cost volume serves as an effective alternative to warping. The result suggests that the deformable cost volume is a suitable replacement of warping and standard cost volume in Devon and potentially other models.

	FlyingChairs Valid	Sintel clean Train
Deformable cost volume	<b>1.87</b>	2.99
Warping	1.88	<b>2.98</b>

Table 3: End point error

	Forward	Backward
Deformable cost volume	<b>50.51</b>	<b>177.17</b>
Warping	57.75	182.75

Table 4: Runtime (ms).

## 6. Discussions

While Devon achieves better results in handling small objects moving fast in the visualization results, it does not outperform multi-resolution based methods such as PWC-Net on Sintel final pass and KITTI. We conjecture that this is due to the fact that Sintel and KITTI mostly contain large objects (e.g. human bodies, cars and buildings), for which the multi-resolution approach might be more suitable. An interesting extension of our work is to combine multi-resolution approach and Devon to handle objects of diverse sizes and speed.

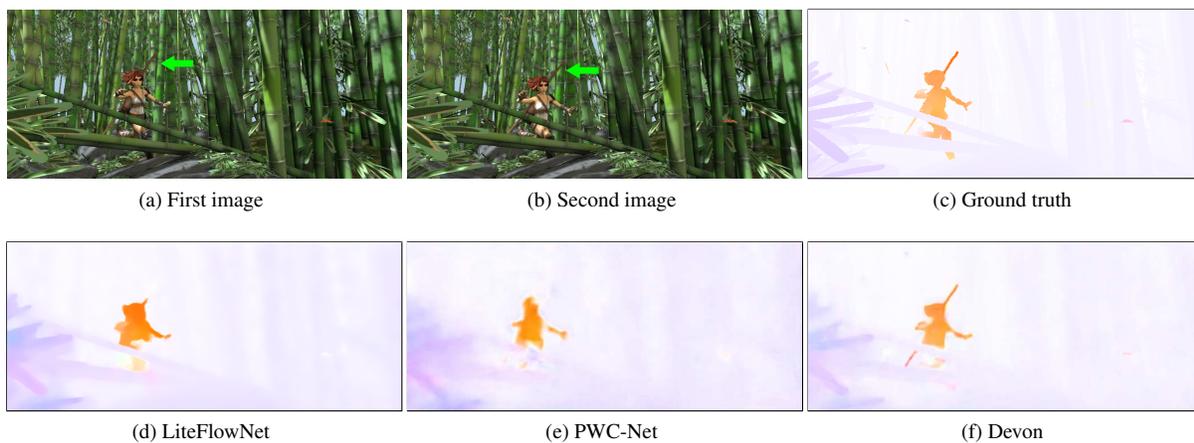


Figure 8: Sintel (training set). Green arrows indicate the small object that moves fast.

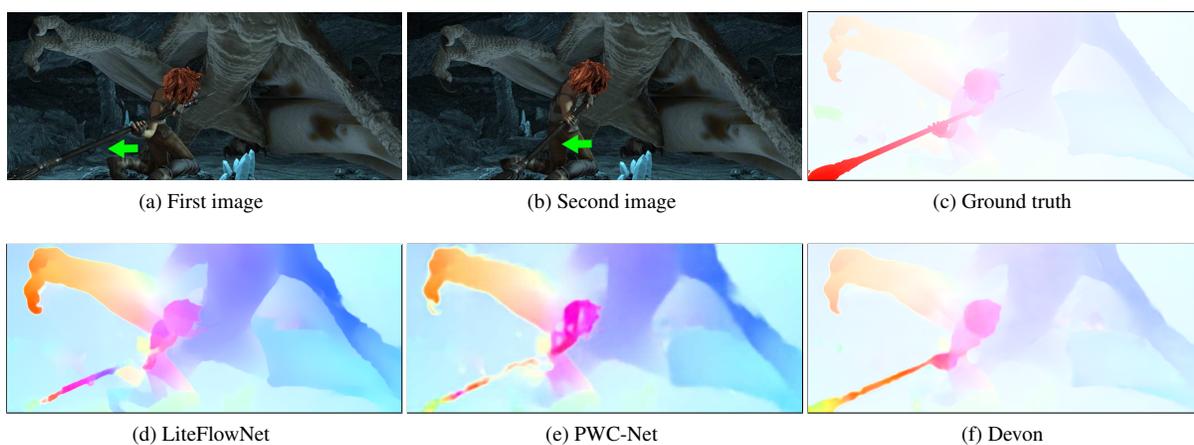


Figure 9: Sintel (training set). Green arrows indicate the small object that moves fast.

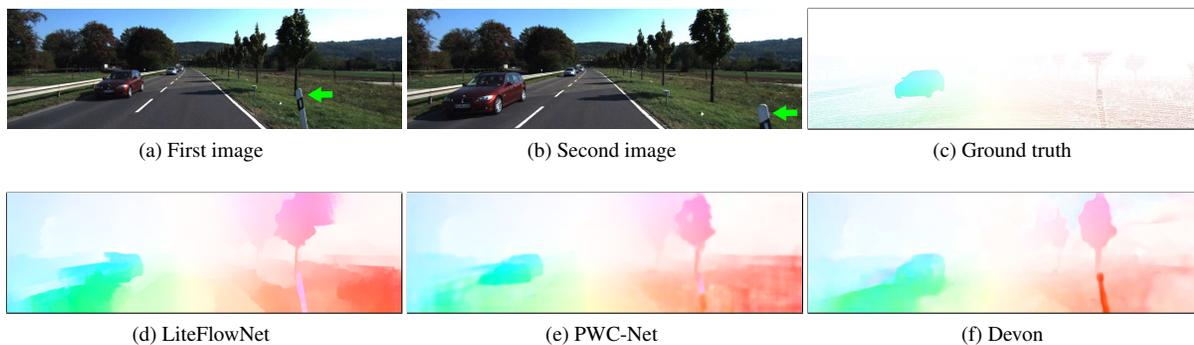


Figure 10: KITTI 2015 (training set). Green arrows indicate the small object that moves fast.

## 7. Related Work

The deformation in deformable cost volume is different from the one in deformable convolutional networks [6]. In deformable cost volume, the cost volume is offset by an external optical flow and dilation. There is no learnable parameter while in deformable convolutional networks, the deformation is element-wise and the offset parameters are learned during training. Another related work is the deformable spatial pyramid matching [22, 16] which uses deformation in the classic energy minimization framework for hierarchical dense matching. Applying normalization on hidden unit outputs is found advantageous in modeling general image relations [25].

The use of dilation in standard cost volume has been proposed in [7] and used in [15], though it has not been used in a multi-scale fashion in each stage as ours to handle the small objects moving fast problem.

Although in this paper we also explored supervised learning of optical flow, the proposed model can be also trained with unsupervised learning losses as in [32, 29, 42, 24].

## 8. Conclusions

In this paper, we proposed a new neural network module, Deformable Cost Volume, which allows the representation of multi-scale motion in a single high resolution and avoids the drawback of warping. Based on it, we designed the Deformable Volume Network, which has three stages estimation all in a single high resolution. The network is demonstrated to be effective in estimating optical flow, especially in situations where small objects move fast. In the future, we expect to combine the strength of our model and multi-resolution models to handle objects of various size and speed in complex scenes. Another key challenge in optical flow estimation is occlusion handling [30, 19, 17]. At the current stage, the cost volume by itself cannot reflect the presence of occlusion. When the matching cost is high, it is difficult to infer whether it is caused by mismatch or occlusion. In the future, we look forward to exploring joint occlusion handling and cost volume construction.

## References

- [1] S. S. Beauchemin and J. L. Barron. The computation of optical flow. *ACM computing surveys*, 1995.
- [2] T. Brox, C. Bregler, and J. Malik. Large displacement optical flow. *CVPR*, 2009.
- [3] T. Brox, A. Bruhn, N. Papenberg, and J. Weickert. High accuracy optical flow estimation based on a theory for warping. *ECCV*, 2004.
- [4] D. J. Butler, J. Wulff, G. B. Stanley, and M. J. Black. A naturalistic open source movie for optical flow evaluation. *ECCV*, 2012.
- [5] L.-C. Chen, G. Papandreou, I. Kokkinos, K. Murphy, and A. L. Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *arXiv*, 2016.
- [6] J. Dai, H. Qi, Y. Xiong, Y. Li, G. Zhang, H. Hu, and Y. Wei. Deformable convolutional networks. *ICCV*, 2017.
- [7] A. Dosovitskiy, P. Fischer, E. Ilg, P. Häusser, C. Hazırbaş, V. Golkov, P. van der Smagt, D. Cremers, and T. Brox. Flownet: Learning optical flow with convolutional networks. *ICCV*, 2015.
- [8] D. Fleet and Y. Weiss. Optical flow estimation. In *Handbook of mathematical models in computer vision*. 2006.
- [9] D. Fortun, P. Bouthemy, and C. Kervrann. Optical flow modeling and computation: a survey. *CVIU*, 2015.
- [10] A. Geiger, P. Lenz, and R. Urtasun. Are we ready for autonomous driving? the kitti vision benchmark suite. *CVPR*, 2012.
- [11] M. Haag and H.-H. Nagel. Combination of edge element and optical flow estimates for 3d-model-based vehicle tracking in traffic image sequences. *IJCV*, 1999.
- [12] K. He, X. Zhang, S. Ren, and J. Sun. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. *CVPR*, 2015.
- [13] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. *CVPR*, 2016.
- [14] B. K. Horn and B. G. Schunck. Determining optical flow. *Artificial Intelligence*, 1981.
- [15] T.-W. Hui, X. Tang, and C. C. Loy. Liteflownet: A lightweight convolutional neural network for optical flow estimation. *CVPR*, 2018.
- [16] J. Hur, H. Lim, C. Park, and S. Chul Ahn. Generalized deformable spatial pyramid: Geometry-preserving dense correspondence estimation. *CVPR*, 2015.
- [17] J. Hur and S. Roth. Iterative residual refinement for joint optical flow and occlusion estimation. 2019.
- [18] E. Ilg, N. Mayer, T. Saikia, M. Keuper, A. Dosovitskiy, and T. Brox. Flownet 2.0: Evolution of optical flow estimation with deep networks. *CVPR*, 2017.
- [19] E. Ilg, T. Saikia, M. Keuper, and T. Brox. Occlusions, motion and depth boundaries with a generic network for disparity, optical flow or scene flow estimation. *ECCV*, 2018.
- [20] M. Jaderberg, K. Simonyan, A. Zisserman, et al. Spatial transformer networks. *NIPS*, 2015.
- [21] J. Janai, F. Güney, A. Ranjan, M. Black, and A. Geiger. Unsupervised learning of multi-frame optical flow with occlusions. *ECCV*, 2018.
- [22] J. Kim, C. Liu, F. Sha, and K. Grauman. Deformable spatial pyramid matching for fast dense correspondences. *CVPR*, 2013.
- [23] D. Kingma and J. Ba. Adam: A method for stochastic optimization. *ICLR*, 2014.
- [24] P. Liu, M. Lyu, I. King, and J. Xu. Selfflow: Self-supervised learning of optical flow. *CVPR*, 2019.
- [25] Y. Lu, Z. Yang, J. Kannala, and S. Kaski. Learning image relations with contrast association networks. *IJCNN*, 2019.
- [26] B. D. Lucas and T. Kanade. An iterative image registration technique with an application to stereo vision. *IJCAI*, 1981.

- [27] N. Mayer, E. Ilg, P. Fischer, C. Hazirbas, D. Cremers, A. Dosovitskiy, and T. Brox. What makes good synthetic training data for learning disparity and optical flow estimation? *IJCV*, 2018.
- [28] N. Mayer, E. Ilg, P. Hausser, P. Fischer, D. Cremers, A. Dosovitskiy, and T. Brox. A large dataset to train convolutional networks for disparity, optical flow, and scene flow estimation. *CVPR*, 2016.
- [29] S. Meister, J. Hur, and S. Roth. Unflow: Unsupervised learning of optical flow with a bidirectional census loss. *AAAI*, 2018.
- [30] M. Neoral, J. Šochman, and J. Matas. Continual occlusion and optical flow estimation. *ACCV*, 2018.
- [31] A. Ranjan and M. J. Black. Optical flow estimation using a spatial pyramid network. *CVPR*, 2017.
- [32] Z. Ren, J. Yan, B. Ni, B. Liu, X. Yang, and H. Zha. Unsupervised deep learning for optical flow estimation. *AAAI*, 2017.
- [33] O. Ronneberger, P. Fischer, and T. Brox. U-net: Convolutional networks for biomedical image segmentation. *MICCAI*, 2015.
- [34] S. Roth and M. J. Black. On the spatial statistics of optical flow. *IJCV*, 2007.
- [35] D. Scharstein and R. Szeliski. A taxonomy and evaluation of dense two-frame stereo correspondence algorithms. *IJCV*, 2002.
- [36] L. Sevilla-Lara, D. Sun, E. G. Learned-Miller, and M. J. Black. Optical flow estimation with channel constancy. *ECCV*, 2014.
- [37] D. Sun, X. Yang, M.-Y. Liu, and J. Kautz. Pwc-net: Cnns for optical flow using pyramid, warping, and cost volume. *CVPR*, 2018.
- [38] D. Teney and M. Hebert. Learning to extract motion from videos in convolutional neural networks. *arXiv*, 2016.
- [39] J. Thewlis, S. Zheng, P. H. Torr, and A. Vedaldi. Fully-trainable deep matching. *BMVC*, 2016.
- [40] B. Ummenhofer, H. Zhou, J. Uhrig, N. Mayer, E. Ilg, A. Dosovitskiy, and T. Brox. Demon: Depth and motion network for learning monocular stereo. *CVPR*, 2017.
- [41] H. Wang, A. Kläser, C. Schmid, and L. Cheng-Lin. Action recognition by dense trajectories. *CVPR*, 2011.
- [42] Y. Wang, Y. Yang, Z. Yang, L. Zhao, P. Wang, and W. Xu. Occlusion aware unsupervised learning of optical flow. 2018.
- [43] J. Xu, R. Ranftl, and V. Koltun. Accurate Optical Flow via Direct Cost Volume Processing. *CVPR*, 2017.
- [44] Y. Yang and S. Soatto. S2f: Slow-to-fast interpolator flow. *CVPR*, 2017.
- [45] F. Yu and V. Koltun. Multi-scale context aggregation by dilated convolutions. *ICLR*, 2015.
- [46] J. Zbontar and Y. LeCun. Stereo matching by training a convolutional neural network to compare image patches. *JMLR*, 2016.