Supplementary Material: Attacking Optical Flow

Anurag Ranjan[†] Joel Janai[†‡] Andreas Geiger[†‡] Michael J. Black[†]
[†]Max Planck Institute for Intelligent Systems
[‡]University of Tübingen
{aranjan, jjanai, ageiger, black}@tue.mpg.de

1. Appendix

This **supplementary document** provides additional results on White-box and Black-box attacks as well as an analysis of FlowNet2 [3] and Back2Future [4] under the Zero-Flow test. In the **video** [1], we show real world attacks using a printed patch placed in the environment.

1.1. White-box Attacks

Additional qualitative results for White-box attacks using patches of size 51×51 and 102×102 are shown in Figure 1 and Figure 2, respectively. We observe that the effect of the patch is more prominent with larger patch sizes. In agreement with the main paper, we note that spatial pyramid architectures are more robust, as compared to encoder-decoder architectures.

1.2. Black-box Attacks

The universal patch is shown in Figure 3. Table 1 shows the performance of optical flow methods when the adversarial patch has zero motion w.r.t. the camera. In comparison to the moving Black-box attacks considered in the main paper, we observe similar effects on all networks and baselines with the adversarial patch. While encoder-decoder networks are strongly affected by the attacks, spatial pyramid networks and classical methods are more robust.

In Figures 4 - 10 we show some additional qualitative results for the Black-box attack with patches moving according to the scene as described in the main paper. These examples demonstrate the feasibility of such attackes in the real world. In Figure 5, for instance, the patch is attached to and moves with a traffic sign, while Figures 6, 9 illustrate cases when a patch is printed on a wall and a car.

	Unattacked	Attacked	
	EPE	EPE	Rel
FlowNet2 [3]	11.90	30.99	+160 %
PWCNet [7]	11.03	11.16	+1 %
FlowNetC [2]	14.56	77.78	+434 %
SpyNet [5]	20.26	20.65	+2 %
Back2Future [4]	17.49	17.76	+2 %
Epic Flow [6]	4.52	4.57	+1 %
LDOF [1]	9.20	9.30	+1 %

Table 1. **Black-box Attacks.** Attacks on different optical flow methods using a universal patch that is static w.r.t. the camera. Methods below the line were not used for training the patch.

	Unattacked		Attacked	
	W Patch	W/O Patch	W Patch	W/O Patch
FlowNetC	14.56	14.56	86.12	80.69
PWCNet	11.03	11.03	11.01	11.08
FlowNet2	11.90	11.90	36.13	34.18
SpyNet	20.26	20.26	20.39	20.50
Back2Future	17.49	17.49	17.44	17.59

Table 2. **Black-box Attacks.** Comparison of the evaluation results with and without considering the attacking patch region.

Evaluation without considering the Patch Region. We also evaulated the effect of the patch without considering the patch region. In case of Black-box attacks (Table 2) the flow outside of the patch region has a similar level of degradation as our results considering the patch region. The unattacked results only show minimal changes below the second decimal place because of the small patch size ($\approx$ 1%).

1.3. Zero-Flow Test

We show feature map visualizations for FlowNet2 and Back2Future under the Zero-Flow test in Figures 11 and 12 respectively. We note that the feature maps of FlowNet2 are not spatially invariant, which is consistant with other networks examined in Section 5 of the main paper. The stacked FlowNetS (part of FlowNet2) seems to be less vulnerable to the adversarial patch as compared to FlowNetC (part of FlowNet2). We also observe that the fusion part of FlowNet2 dramatically amplifies the degradations in optical flow predictions. The deconvolution layers show similar checkerboard artifacts as FlowNetC and PWC-Net analysed in Section 5 of the main paper.

[1] http://flowattack.is.tue.mpg.de/

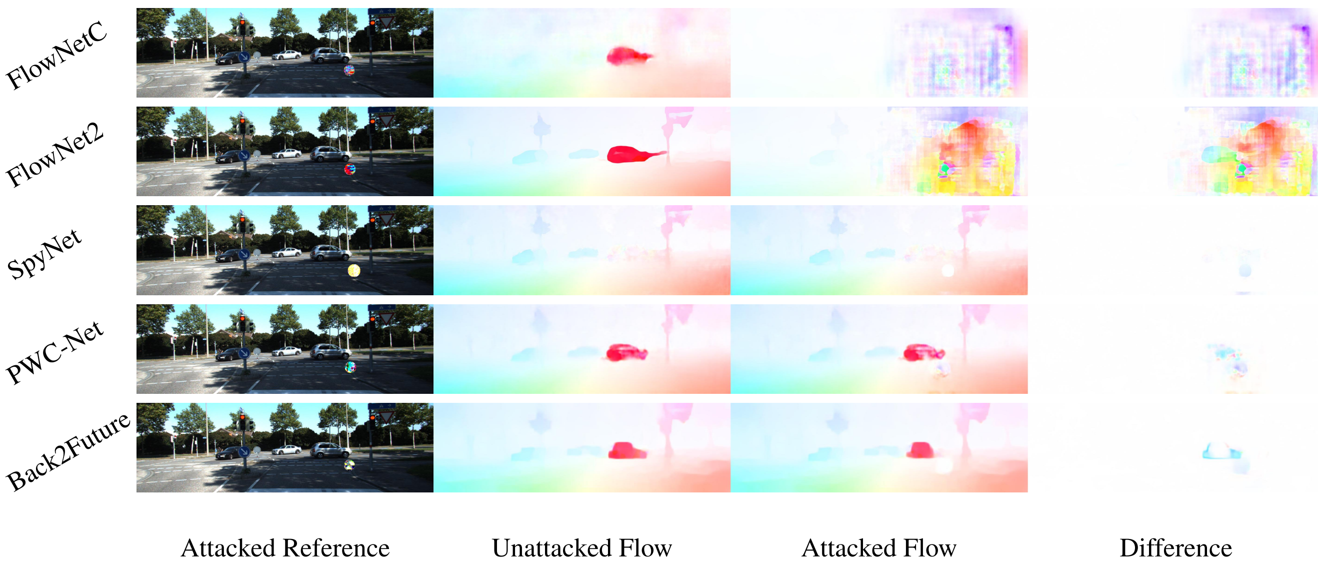

Figure 1. **White-box Attacks** on all networks using 51x51 patches.

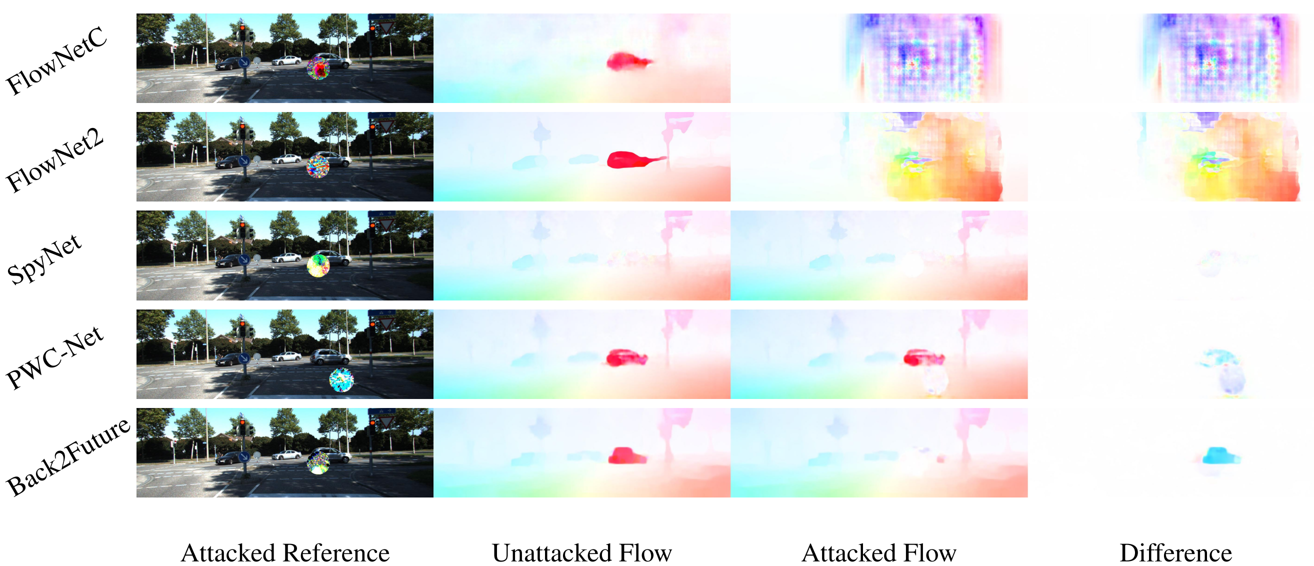

Figure 2. **White-box Attacks** on all networks using 102x102 patches.

For Back2Future, we note that, although the feature maps are not spatially invariant, their magnitude remains small irrespective of the presence or absence of the adversary. Interestingly, Back2Future gives reasonable flow predictions at coarser levels of the pyramid unlike PWC-Net, even though they share a common architecture.

We note that the problem of spatially variant feature maps continue across all the examined networks, along with the checkerboard artifacts.

References

[1] T. Brox and J. Malik. Large displacement optical flow: Descriptor matching in variational motion estimation. *IEEE Trans. on Pattern Analysis and Machine Intelligence (PAMI)*, 2011. 1

[2] A. Dosovitskiy, P. Fischery, E. Ilg, C. Hazirbas, V. Golkov, P. van der Smagt, D. Cremers, T. Brox, et al. Flownet: Learning optical flow with convolutional networks. In *Proc. of the IEEE International Conf. on Computer Vision (ICCV)*, pages 2758–2766. IEEE, 2015. 1

[3] E. Ilg, N. Mayer, T. Saikia, M. Keuper, A. Dosovitskiy, and T. Brox. Flownet 2.0: Evolution of optical flow estimation

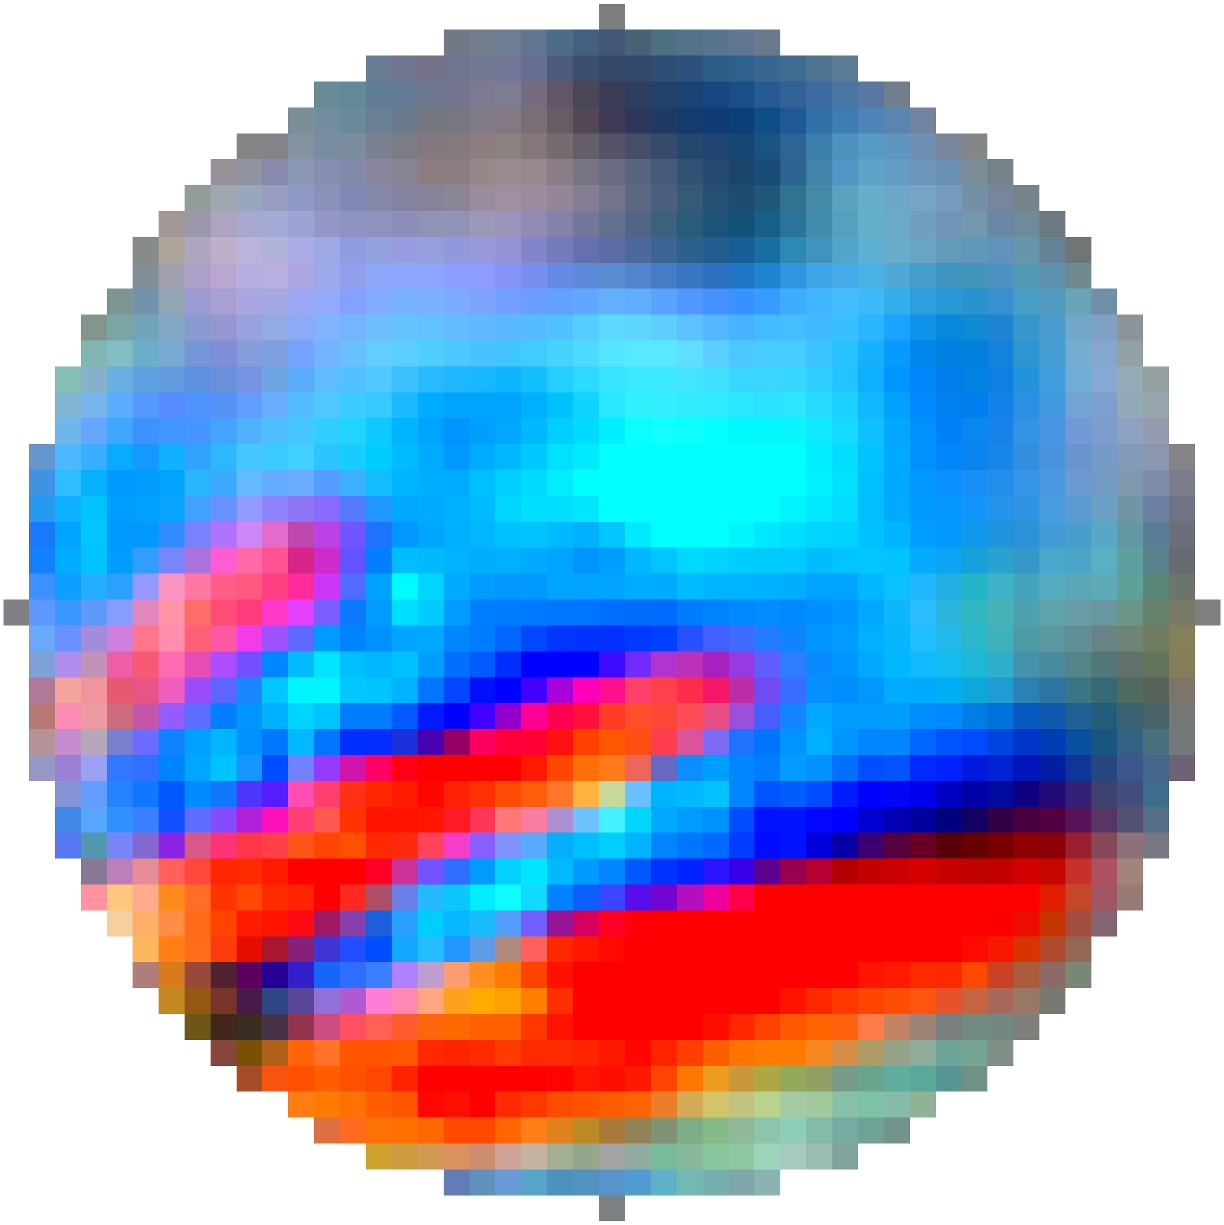

Figure 3. **"Universal" Patch** obtained by optimizing over FlowNet2 and PWCNet. Patch is enlarged for visualization.

with deep networks. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, volume 2, 2017. 1

[4] J. Janai, F. Güney, A. Ranjan, M. Black, and A. Geiger. Unsupervised learning of multi-frame optical flow with occlusions. In *Proc. of the European Conf. on Computer Vision (ECCV)*, 2018. 1

[5] A. Ranjan and M. J. Black. Optical flow estimation using a spatial pyramid network. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2017. 1

[6] J. Revaud, P. Weinzaepfel, Z. Harchaoui, and C. Schmid. EpicFlow: Edge-preserving interpolation of correspondences for optical flow. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2015. 1

[7] D. Sun, X. Yang, M. Liu, and J. Kautz. Pwc-net: Cnns for optical flow using pyramid, warping, and cost volume. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 8934–8943, 2018. 1

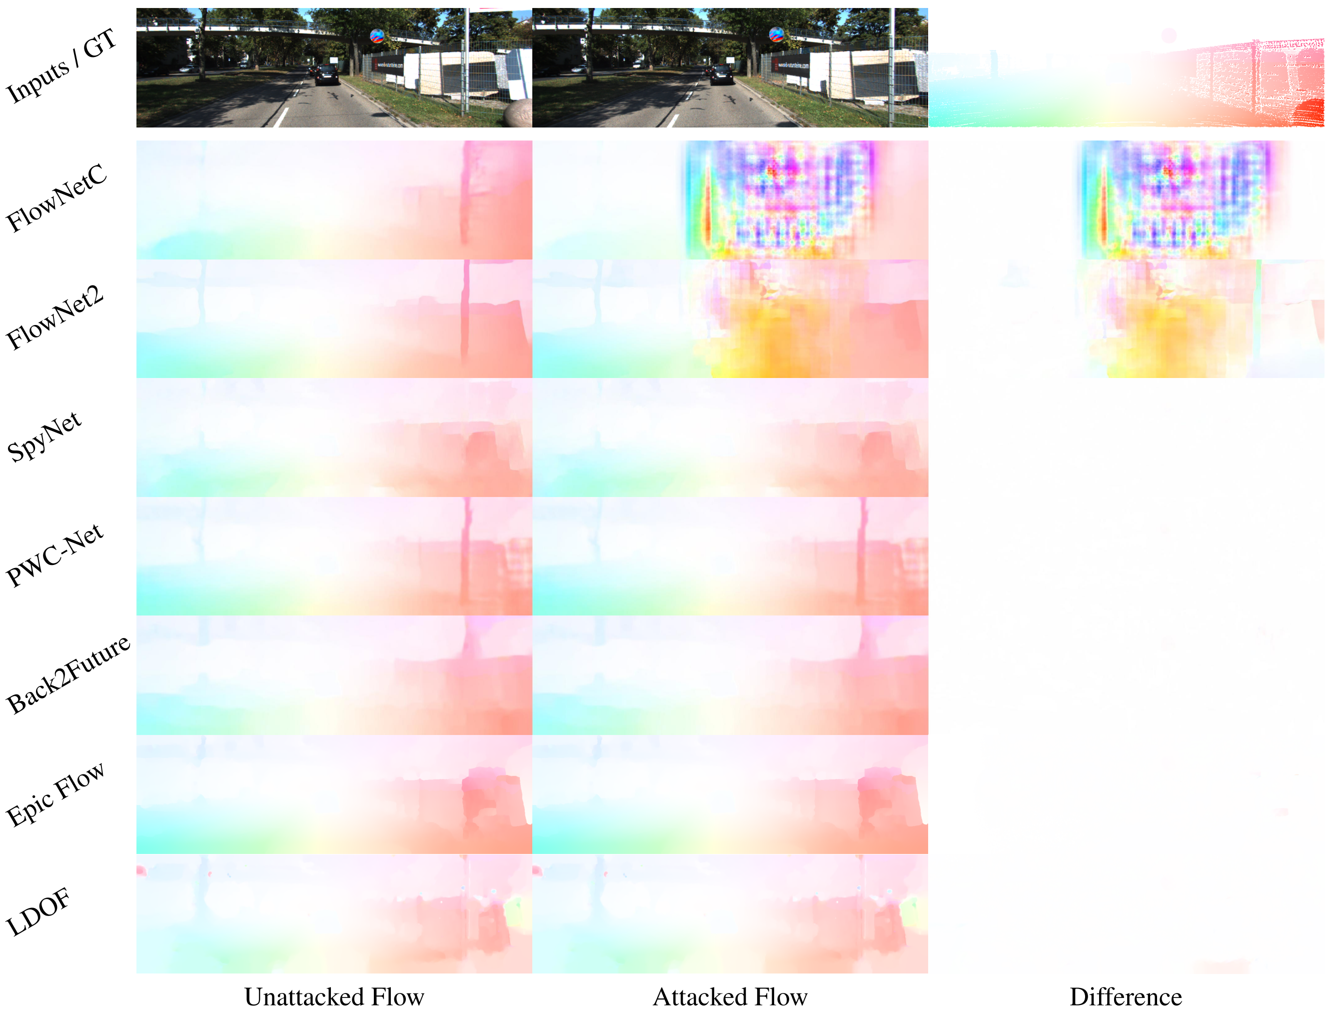

Figure 4. **Black-box Attacks.** Universal patch trained on FlowNet2 and PWC-Net used on all approaches. For this evaluation, we move the patch according to the static scene.

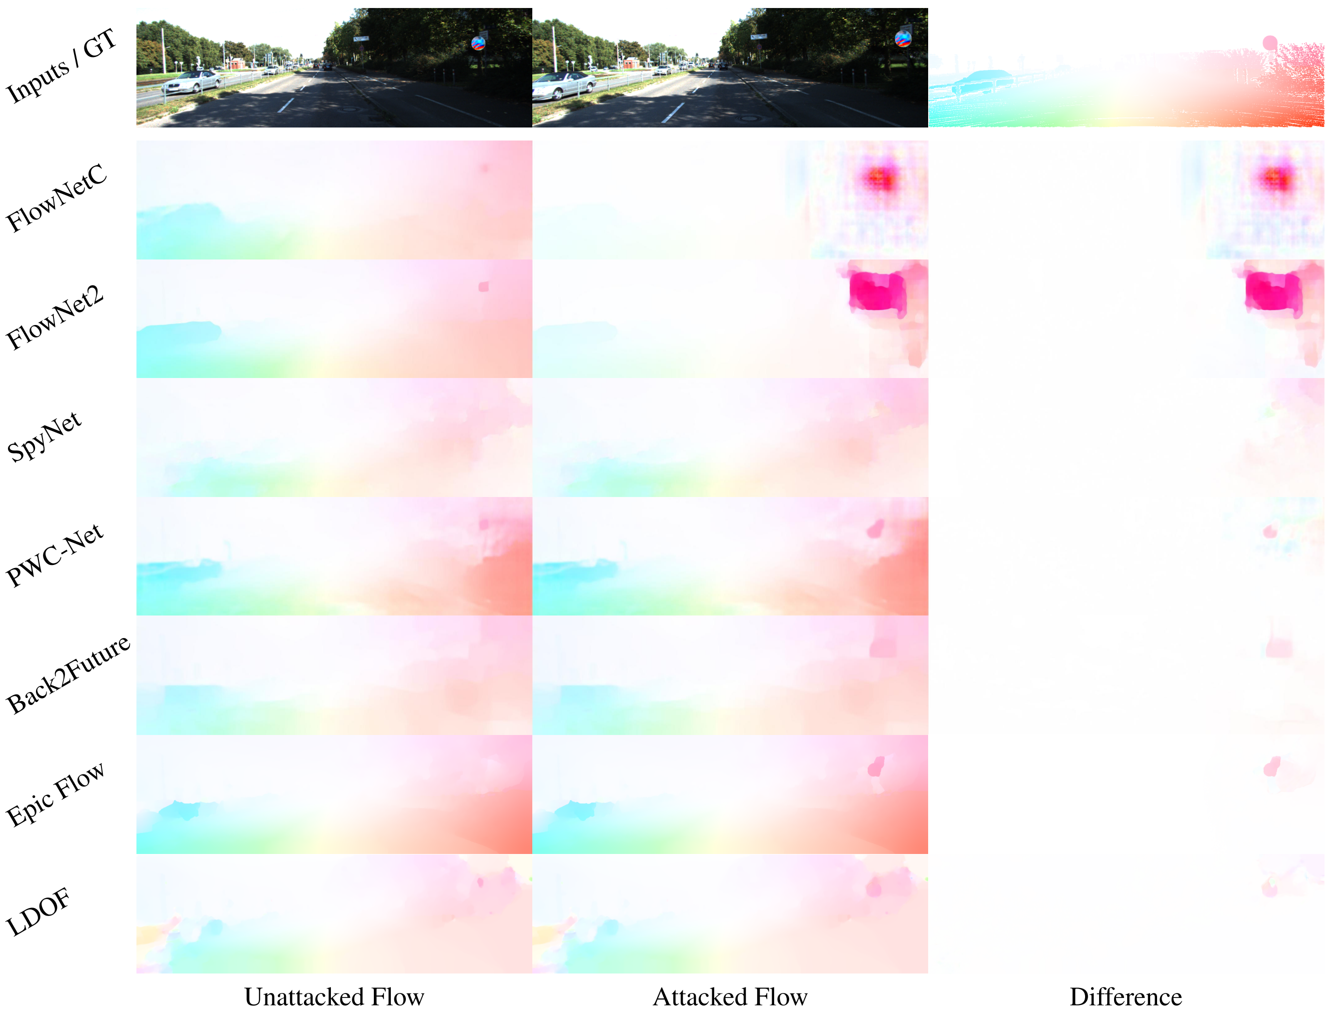

Figure 5. **Black-box Attacks.** Universal patch trained on FlowNet2 and PWC-Net used on all approaches. For this evaluation, we move the patch according to the static scene.

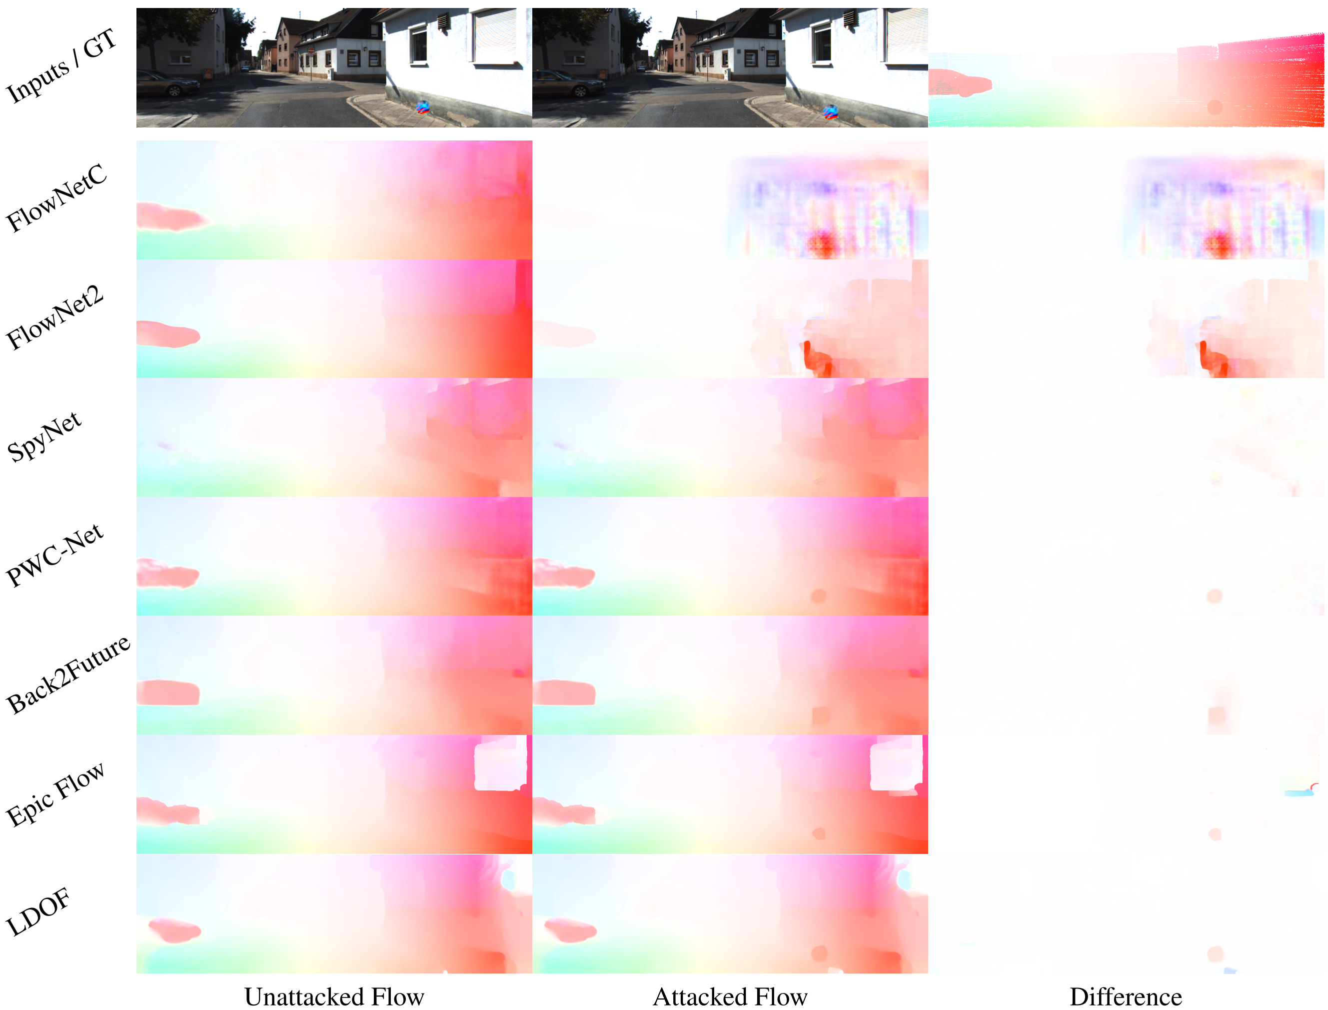

Figure 6. **Black-box Attacks.** Universal patch trained on FlowNet2 and PWC-Net used on all approaches. For this evaluation, we move the patch according to the static scene.

Figure 7. **Black-box Attacks.** Universal patch trained on FlowNet2 and PWC-Net used on all approaches. For this evaluation, we move the patch according to the static scene.

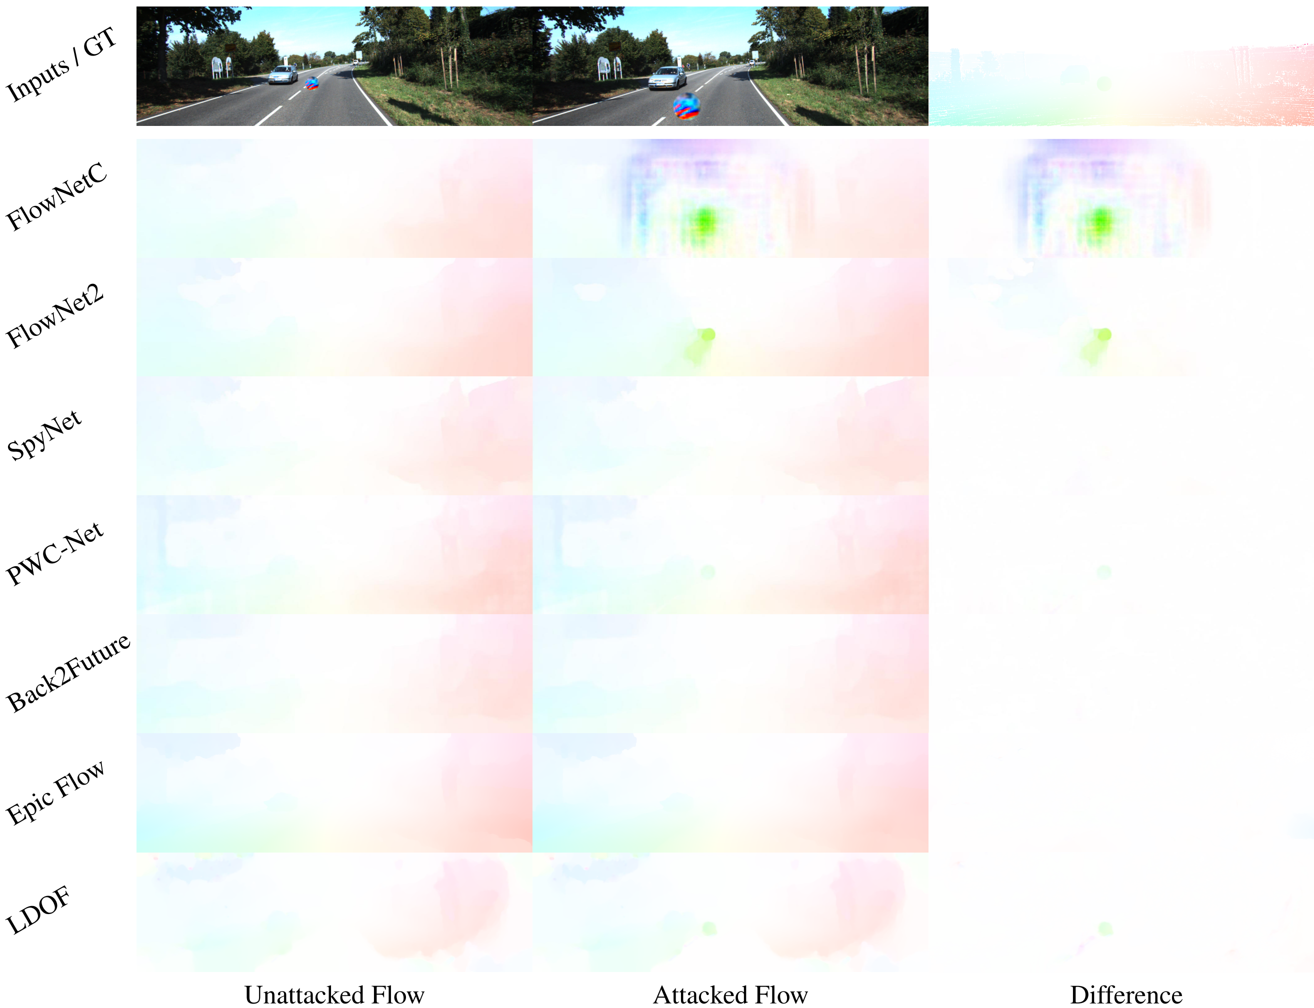

Figure 8. **Black-box Attacks.** Universal patch trained on FlowNet2 and PWC-Net used on all approaches. For this evaluation, we move the patch according to the static scene.

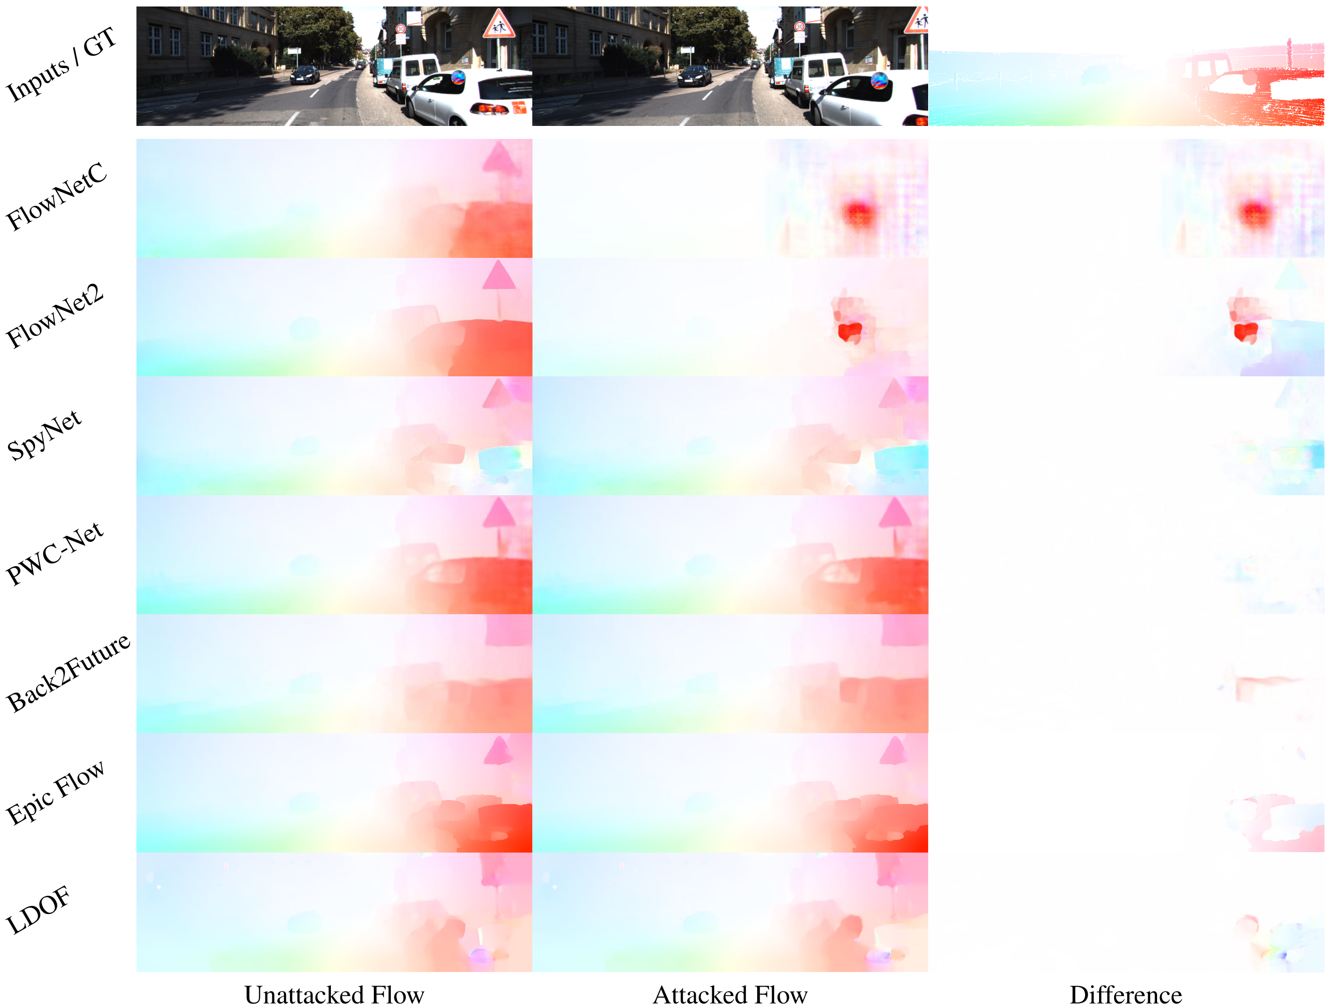

Figure 9. **Black-box Attacks.** Universal patch trained on FlowNet2 and PWC-Net used on all approaches. For this evaluation, we move the patch according to the static scene.

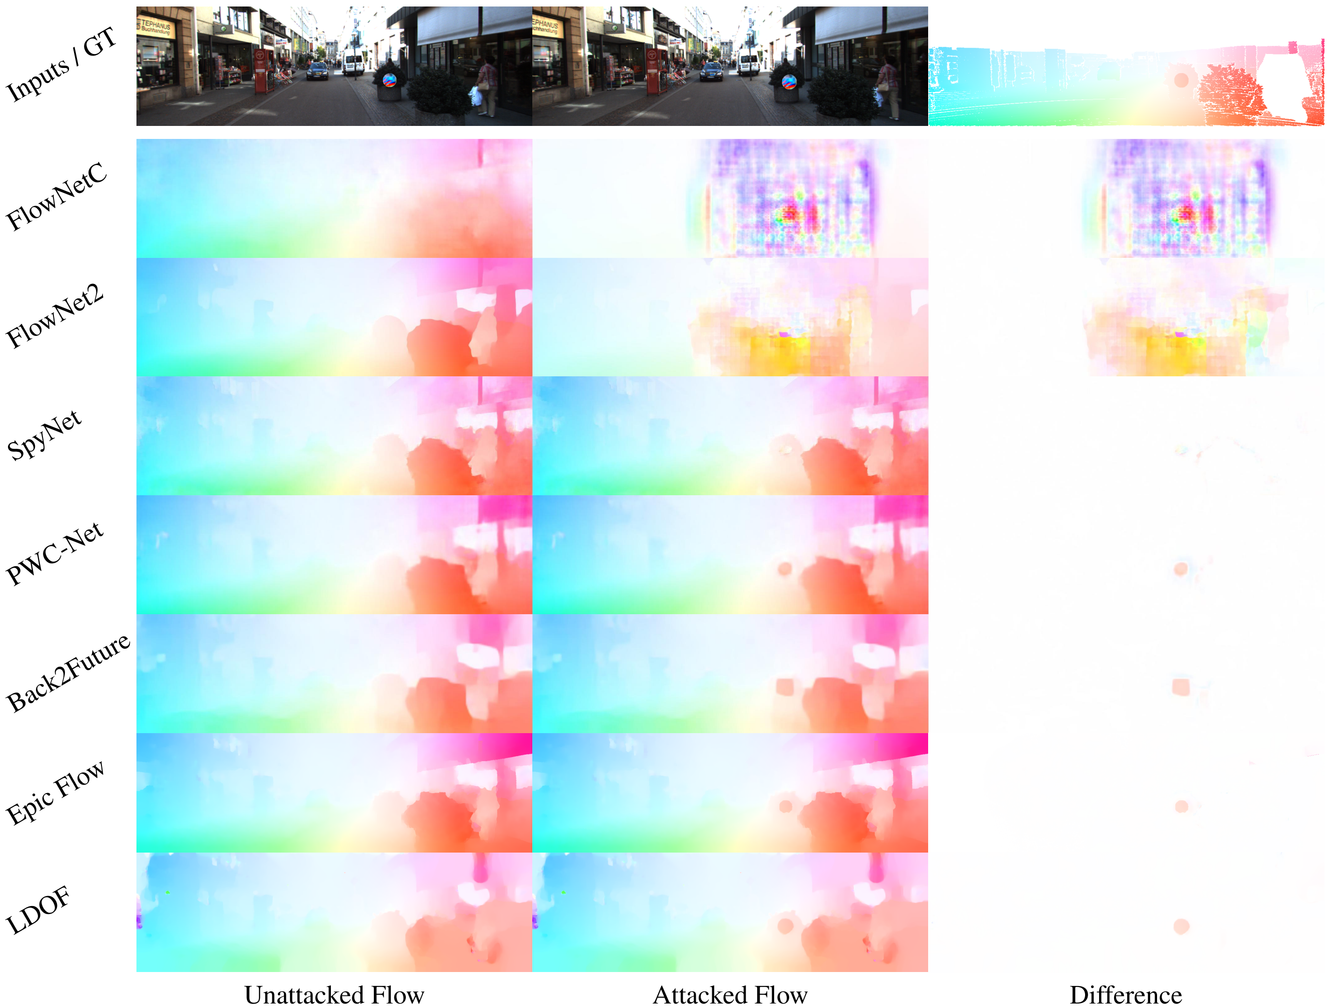

Figure 10. **Black-box Attacks.** Universal patch trained on FlowNet2 and PWC-Net used on all approaches. For this evaluation, we move the patch according to the static scene.

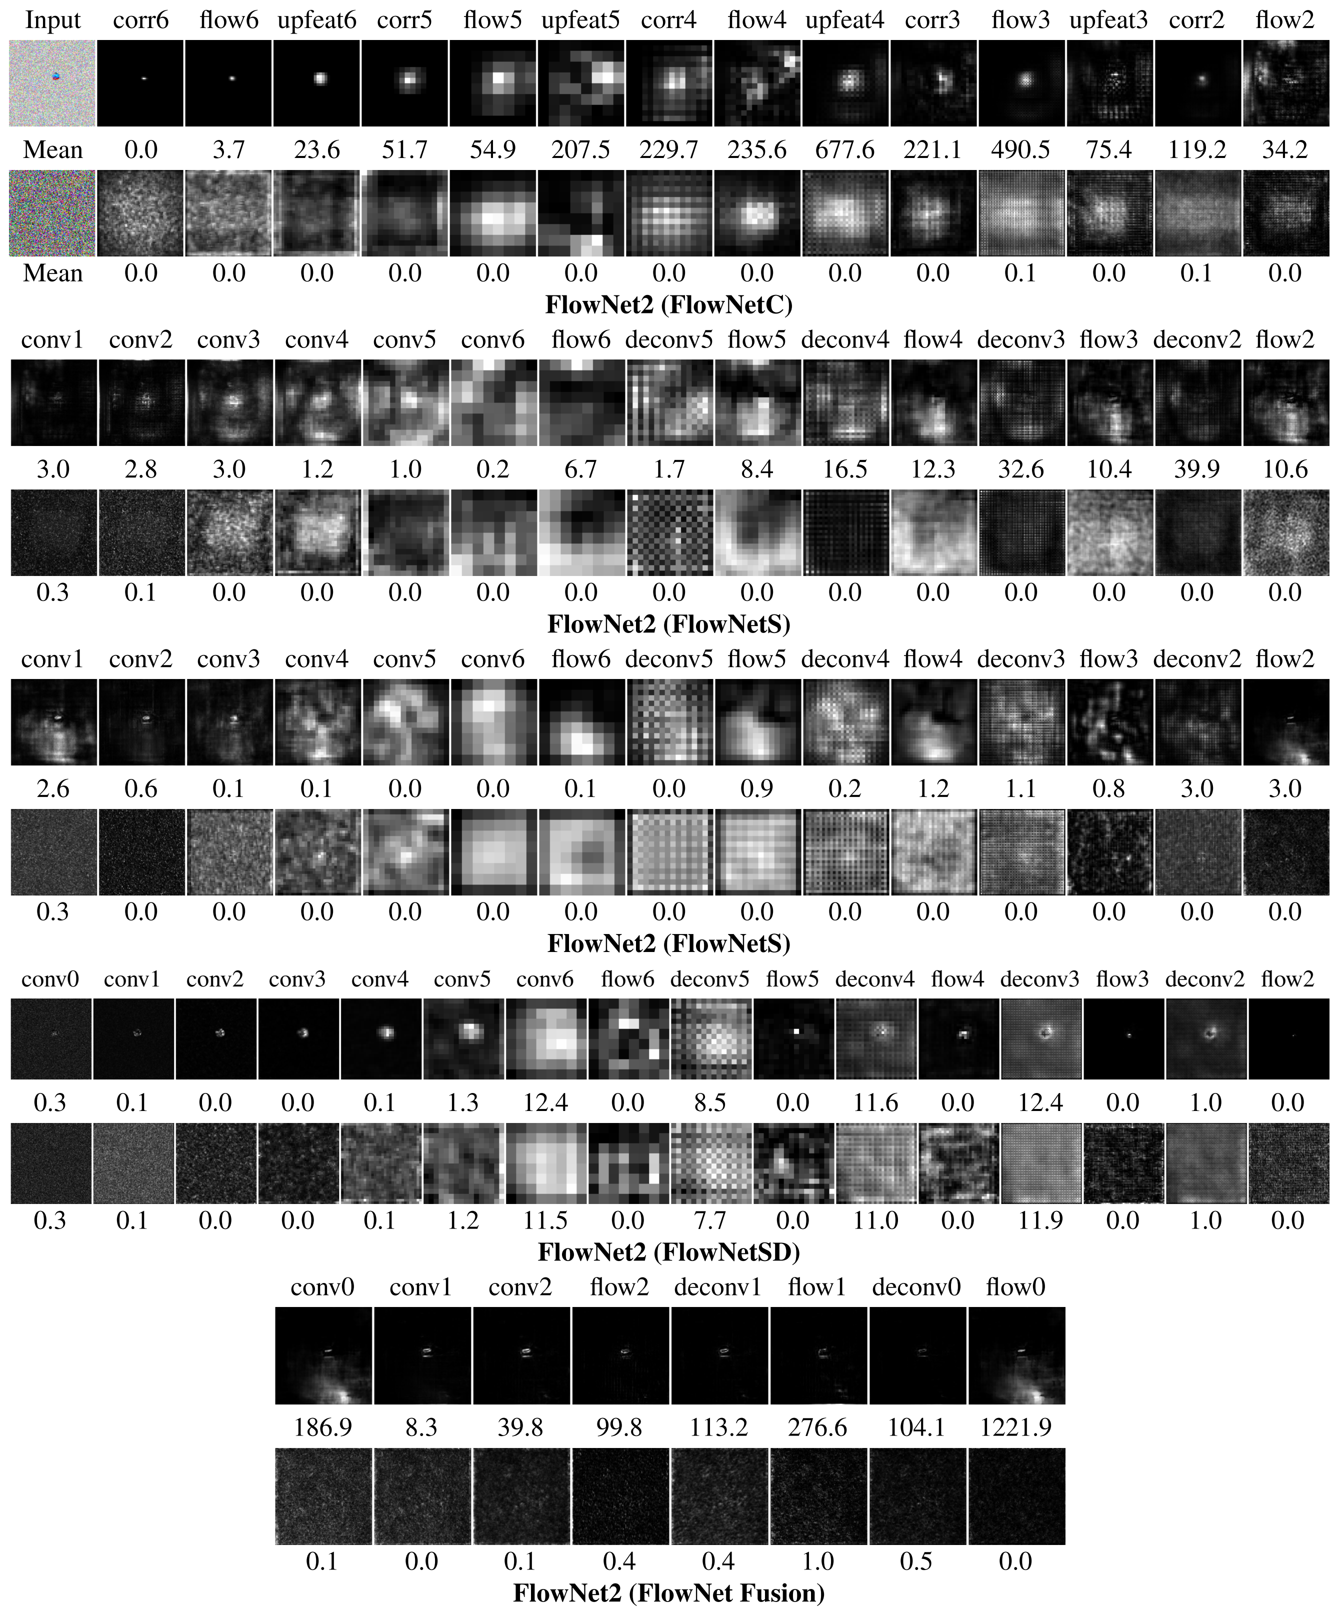

Figure 11. **Zero-Flow Test.** Feature maps of Flownet2 under the Zero-Flow test. Top to bottom, we show rows corresponding to FlowNetC, FlowNetS, FlowNetS, FlowNetSD and FlowNet Fusion that constitute FlowNet2.

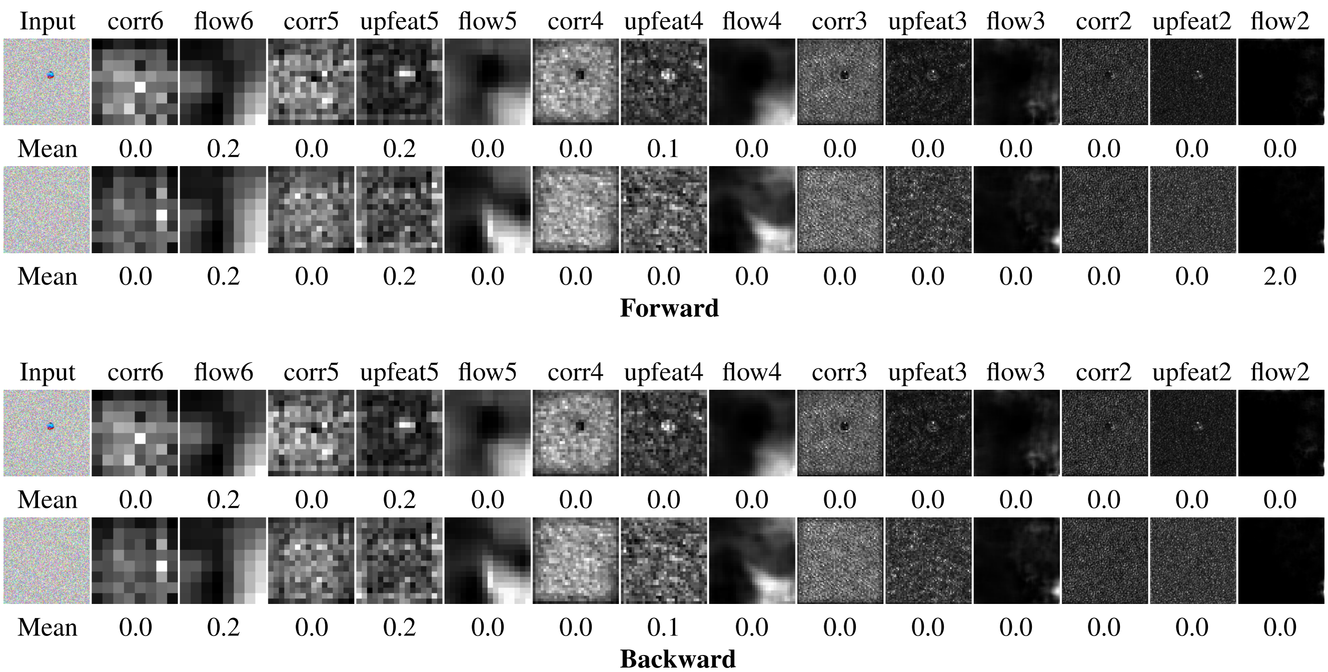

Figure 12. **Zero-Flow Test.** Feature maps of Back2Future under Zero-Flow test. Top to bottom, we show rows corresponding to forward and backward parts of Back2Future in a multi-frame set up.