Distracting Downpour: Adversarial Weather Attacks for Motion Estimation –Supplementary Material–

Jenny Schmalfuss Lukas Mehl Andrés Bruhn
Institute for Visualization and Interactive Systems, University of Stuttgart
firstname.lastname@vis.uni-stuttgart.de

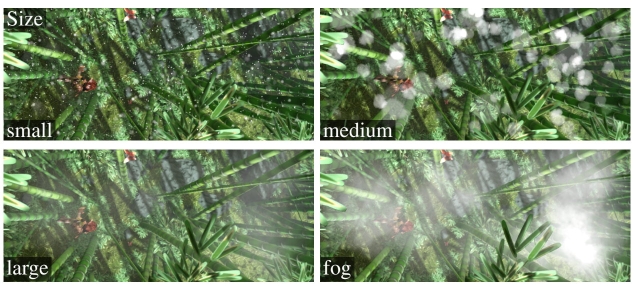

Figure A.1. Augmentations for different *particle sizes* and *transparencies* from Main Tab. 2 on an exemplary Sintel [1] frame. Augmentations for *particle count*, *motion blur* and *color* are shown in Fig. A.2, augmentation parameters are listed in Tab. A.1.

A. Additional Material for Experiments

We provide the code to generate the weather augmentations and run all adversarial weather attacks at https://github.com/cv-stuttgart/DistractingDownpour. The tested optical flow networks utilize the respective author-provided PyTorch implementations with Sintel-checkpoints for FlowFormer [2], GMA [10], RAFT [10] and FlowNetCRobust [9]. For FlowNet2 [3] and SpyNet [6] we use the implementations from [7] and [5], respectively.

A.1. Weather augmentations: Configurations

In Tab. A.1 we give a full list of parameter configurations for the particle effects from Main Tab. 2. In addition to the weather visualizations in Main Fig. 4, we visualize all *Size*-variations for particles in Fig. A.1, and all *Particle count*, *Motion blur* and *color* variations in Fig. A.2. From these figures it becomes clear that the configurations *size: small*, *motion blur: 0.0* and *color: white* all visually correspond to snow. Therefore, they were not chosen as four main weather effects. Instead, we selected configurations that lead to more diverse visual appearance, even though these configurations were not necessarily the most effective ones to perturb the optical flow output in Main Tab. 2.

A.2. Adversarial weather attacks

A.2.1 Attack Configurations

With the provided code and the network implementations above, Tab. A.2 lists the configurations for all weather attacks that were used to create Tables 4, 5 and 6 from the Main paper. To compare to PCFA [8] and I-FGSM [9], we use the implementation from [8] and the author-provided configurations for PCFA with $\varepsilon_2 = $ 5e-3, AEE loss, COV constraint and disjoint, non-universal perturbations. For I-FGSM we use a perturbation bound of $\varepsilon_\infty = $ 5e-3 and 25 optimization steps. Both attacks are run on Sintel train.

A.2.2 Additional configurations for training with snow

Regarding the configurations for the snow-augmented training in Sec. 4.3, Main Tab. 7 uses snow, rain, sparks and fog augmentations as specified in Tab. A.1 and the respective attack configurations that were used for Main Tab. 5, which are listed in Tab. A.2.

A.2.3 Additional visualizations for weather attacks

Finally, in Figures A.3, A.4, A.5 and A.6 we provide additional visualizations for attacks with snow, rain, sparks and fog. They complement Main Fig. 6, and provide visualizations for sample images from the attack runs in Main Tab. 5.

References

[1] Daniel Butler, Jonas Wulff, Garrett Stanley, and Michael J. Black. A naturalistic open source movie for optical flow evaluation. In *Proc. European Conference on Computer Vision (ECCV)*, pages 611–625, 2012.

[2] Zhaoyang Huang, Xiaoyu Shi, Chao Zhang, Qiang Wang, Ka Chun Cheung, Hongwei Qin, Jifeng Dai, and Hongsheng Li. FlowFormer: A transformer architecture for optical flow. In *Proc. European Conference on Computer Vision (ECCV)*, pages 668–685, 2022.

	Weather	Count	Size	d-decay	Templates	(R,G,B)	(δH, δL, δS)	Mode	θ	m_y	$\delta\angle m$	$\delta\|m\|$	Blur	Length	Particles
		Particle base properties				Color properties				Motion properties			Motion blur		
Particles	1000	1000	71	9	particles	(255,255,255)	(0, 0.0, 0.0)	additive	0.75	0.2	0.0	0	-	0.0	0
	2000	2000	71	9	particles	(255,255,255)	(0, 0.0, 0.0)	additive	0.75	0.2	0.0	0	-	0.0	0
	3000 (snow)	3000	71	9	particles	(255,255,255)	(0, 0.0, 0.0)	additive	0.75	0.2	0.0	0	-	0.0	0
	4000	4000	71	9	particles	(255,255,255)	(0, 0.0, 0.0)	additive	0.75	0.2	0.0	0	-	0.0	0
	5000	5000	71	9	particles	(255,255,255)	(0, 0.0, 0.0)	additive	0.75	0.2	0.0	0	-	0.0	0
	grey	3000	71	9	particles	(127,127,127)	(0, 0.0, 0.0)	Meshkin	0.75	0.2	0.0	0	-	0.0	0
Motion blur	0.0	3000	71	9	particles	(255,255,255)	(0, 0.0, 0.0)	additive	0.75	0.2	0.0	0	-	0.0	0
	0.0375	3000	67	9	particles	(255,255,255)	(0, 0.0, 0.0)	additive	0.75	0.2	0.1	4	✓	0.0375	20
	0.075	3000	61	9	particles	(255,255,255)	(0, 0.0, 0.0)	additive	0.75	0.2	0.1	4	✓	0.075	20
	0.1125	3000	57	9	particles	(255,255,255)	(0, 0.0, 0.0)	additive	0.75	0.2	0.1	4	✓	0.1125	20
	0.15 (rain)	3000	51	9	particles	(255,255,255)	(0, 0.0, 0.0)	additive	0.75	0.2	0.1	4	✓	0.15	20
Color – Additive	white	3000	41	9	particles	(255,255,255)	(0, 0.0, 0.0)	additive	1.5	-0.05	0.2	60	✓	0.3	10
	red (sparks)	3000	41	9	particles	(191, 79, 64)	(15, 0.1, 0.1)	additive	1.5	-0.05	0.2	60	✓	0.3	10
	green	3000	41	9	particles	(79,191, 64)	(15, 0.1, 0.1)	additive	1.5	-0.05	0.2	60	✓	0.3	10
	blue	3000	41	9	particles	(64, 79,191)	(15, 0.1, 0.1)	additive	1.5	-0.05	0.2	60	✓	0.3	10
	color	3000	41	9	particles	(205, 80, 80)	(180, 0.1, 0.1)	additive	1.5	-0.05	0.2	60	✓	0.3	10
Color – α-Blending	white	3000	41	9	particles	(255,255,255)	(0, 0.0, 0.0)	Meshkin	1.5	-0.05	0.2	60	✓	0.3	10
	red	3000	41	9	particles	(191, 79, 64)	(15, 0.1, 0.1)	Meshkin	1.5	-0.05	0.2	60	✓	0.3	10
	green	3000	41	9	particles	(79,191, 64)	(15, 0.1, 0.1)	Meshkin	1.5	-0.05	0.2	60	✓	0.3	10
	blue	3000	41	9	particles	(64, 79,191)	(15, 0.1, 0.1)	Meshkin	1.5	-0.05	0.2	60	✓	0.3	10
	color	3000	41	9	particles	(205, 80, 80)	(180, 0.1, 0.1)	Meshkin	1.5	-0.05	0.2	60	✓	0.3	10
Size	small	3000	71	9	dust	(255,255,255)	(0, 0.0, 0.0)	Meshkin	1.0	0.0	0.0	0	-	0.0	0
	medium	141	161	1.75	dust	(255,255,255)	(0, 0.0, 0.0)	Meshkin	0.78	0.0	0.0	0	-	0.0	0
	large	60	451	0.8	dust	(255,255,255)	(0, 0.0, 0.0)	Meshkin	0.1	0.0	0.0	0	-	0.0	0
	fog (fog)	60	451	0.8	dust	(255,255,255)	(0, 0.0, 0.0)	Meshkin	0.25	0.0	0.0	0	-	0.0	0

Table A.1. *Particle configurations* for Sintel train dataset augmentations from Main Tab. 2. It additionally lists the *grey* particle configuration used in Main Tab. 4. d-Decay is a depth-decay parameter that affects the size, δH, δL and δS are random color variations in the HLS space around the initial RGB color configurations. While all effects use a depth-dependent transparency scaling, fog has a depth-constant transparency of 0.3. The motion m is always along the y-axis (vertically, $m_x = m_z = 0$), and may vary by a random angle $\delta\angle m$ or be scaled by a random factor that scales with a fraction of $\|m\|$. All configurations were created with 8 GPUs and a random seed of 0. To train RAFT on another snow dataset than the test set (Main Tab. 7), the training set uses a random seed of 1234.

Tab.	Dataset	Augment.	LR	δ_{p_1}	δ_{p_2}	δ_γ	δ_θ
Table 4	Sintel-tr115	grey	1e-5	✓	-	-	-
	Sintel-tr115	grey	1e-5	-	✓	-	-
	Sintel-tr115	grey	1e-3	-	-	✓	-
	Sintel-tr115	grey	1e-3	-	-	-	✓
	Sintel-tr115	grey	1e-5	✓	✓	-	-
	Sintel-tr115	grey	1e-3	-	-	✓	✓
	Sintel-tr115	grey	1e-5	✓	✓	✓	✓
Table 5	Sintel-tr115	snow	1e-5	✓	✓	✓	✓
	Sintel-tr115	rain	1e-5	✓	✓	✓	✓
	Sintel-tr115	sparks	1e-5	✓	✓	✓	✓
	Sintel-tr115	fog	1e-5	✓	-	✓	✓
T.6	Sintel train	snow	1e-5	✓	✓	✓	✓

Table A.2. *Weather attack configurations* for the experiments from Main Tables 4, 5 and 6. *Augment* specifies the augmentation, *cf*. Tab. A.1, *LR* denotes the optimizer learning rate and the optimization variables $\delta_{p_1}, \delta_{p_2}, \delta_\gamma$ and δ_θ indicate which of them were optimized. Optimization with 750 steps of Adam using weights $\alpha_1 = \alpha_2 = 1000$ for the loss function.

[3] 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 *Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2462–2470, 2017.

[4] Shihao Jiang, Dylan Campbell, Yao Lu, Hongdong Li, and Richard Hartley. Learning to estimate hidden motions with global motion aggregation. In *Proc. IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 9772–9781, 2021.

[5] Simon Niklaus. A reimplementation of SPyNet using PyTorch, 2018.

[6] Anurag Ranjan and Michael J. Black. Optical flow estimation using a spatial pyramid network. In *Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4161–4170, 2017.

[7] Fitsum Reda, Robert Pottorff, Jon Barker, and Bryan Catanzaro. flownet2-pytorch: Pytorch implementation of FlowNet 2.0: Evolution of optical flow estimation with deep networks, 2017.

[8] Jenny Schmalfuss, Philipp Scholze, and Andrés Bruhn. A perturbation-constrained adversarial attack for evaluating the robustness of optical flow. In *Proc. European Conference on Computer Vision (ECCV)*, pages 183–200, 2022.

[9] Simon Schrodi, Tonmoy Saikia, and Thomas Brox. Towards understanding adversarial robustness of optical flow

Figure A.2. Augmentations for different *particle count*, *motion blur* and *color* from Main Tab. 2. on exemplary Sintel [1] frames. Augmentations for *particle sizes* are shown in Fig. A.1, augmentation parameters are listed in Tab. A.1.

networks. In *Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8916–8924, 2022.

[10] Zachary Teed and Jia Deng. RAFT: Recurrent all-pairs field transforms for optical flow. In *Proc. European Conference on Computer Vision (ECCV)*, pages 402–419, 2020.

Figure A.3. *Snow*. Qualitative results for 3000 *snowflakes* on images from the Sintel final dataset with *random* initialization and *adversarial* optimization with optical flow predictions for FlowNet2 [3], FlowNetCRobust [9], SpyNet [6], RAFT [10], GMA [4] and FlowFormer [2], as extension to Main Fig. 6. See also Figs. A.4, A.5 and A.6.

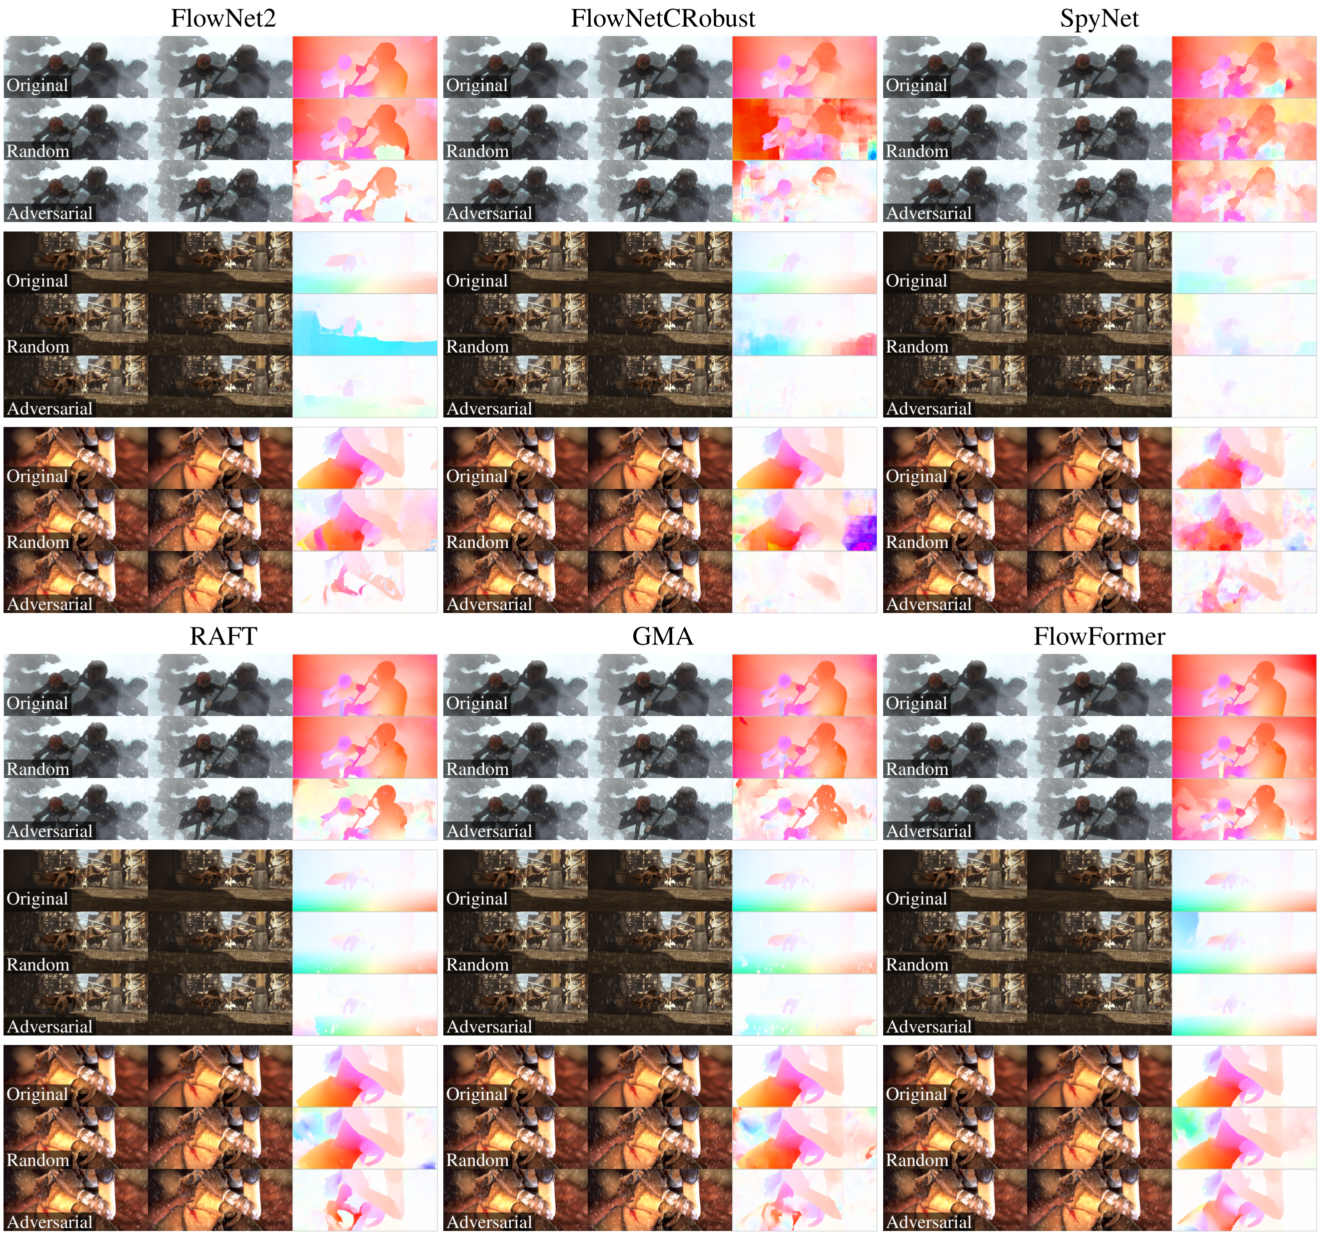

Figure A.4. *Rain*. Qualitative results for 3000 *rain streaks* on images from the Sintel final dataset with *random* initialization and *adversarial* optimization with optical flow predictions for FlowNet2 [3], FlowNetCRobust [9], SpyNet [6], RAFT [10], GMA [4] and FlowFormer [2] as extension to Main Fig. 6. See also Figs. A.3, A.5 and A.6.

Figure A.5. *Sparks*. Qualitative results for 3000 *fire sparks* on images from the Sintel final dataset with *random* initialization and *adversarial* optimization with optical flow predictions for FlowNet2 [3], FlowNetCRobust [9], SpyNet [6], RAFT [10], GMA [4] and FlowFormer [2] as extension to Main Fig. 6 and visualization of exemplary results from Main Tab. 5. See also Figs. A.3, A.4 and A.6.

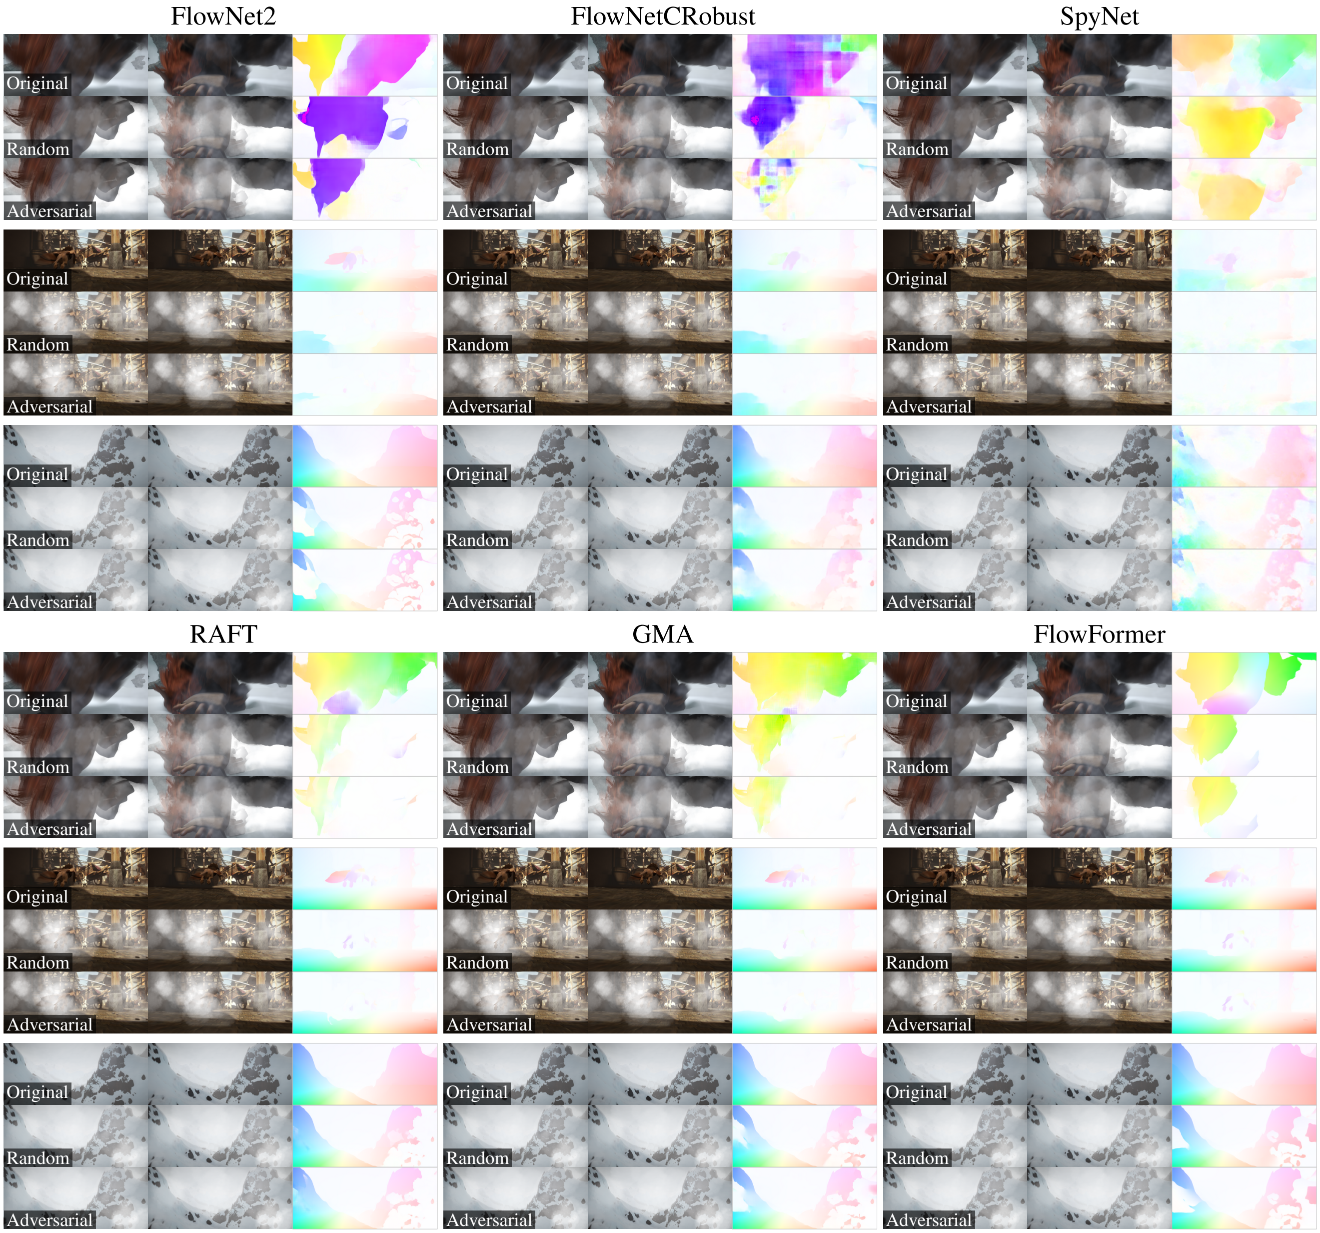

Figure A.6. *Fog.* Qualitative results for 60 *fog clouds* on images from the Sintel final dataset with *random* initialization and *adversarial* optimization with optical flow predictions for FlowNet2 [3], FlowNetCRobust [9], SpyNet [6], RAFT [10], GMA [4] and FlowFormer [2], as extension to Main Fig. 6. See also Figs. A.3, A.4 and A.5.