

Detection Defenses: An Empty Promise against Adversarial Patch Attacks on Optical Flow

Erik Scheurer^{*1} Jenny Schmalfluss^{*2} Alexander Lis Andrés Bruhn²
 Institute for Visualization and Interactive Systems, University of Stuttgart
 first.last@{¹simtech,²vis}.uni-stuttgart.de

Abstract

Adversarial patches undermine the reliability of optical flow predictions when placed in arbitrary scene locations. Therefore, they pose a realistic threat to real-world motion detection and its downstream applications. Potential remedies are defense strategies that detect and remove adversarial patches, but their influence on the underlying motion prediction has not been investigated. In this paper, we thoroughly examine the currently available detect-and-remove defenses ILP and LGS for a wide selection of state-of-the-art optical flow methods, and illuminate their side effects on the quality and robustness of the final flow predictions. In particular, we implement defense-aware attacks to investigate whether current defenses are able to withstand attacks that take the defense mechanism into account. Our experiments yield two surprising results: Detect-and-remove defenses do not only lower the optical flow quality on benign scenes, in doing so, they also harm the robustness under patch attacks for all tested optical flow methods except FlowNetC. As currently employed detect-and-remove defenses fail to deliver the promised adversarial robustness for optical flow, they evoke a false sense of security. The code is available at <https://github.com/cv-stuttgart/DetectionDefenses>

1. Introduction

Adversarial attacks have an enormous potential to mislead optical flow methods into predicting the wrong apparent 2D motion from image sequences. Among them, adversarial patches [35, 47, 50] are the most safety-critical as they distort the optical flow predictions largely independent of their location and orientation, cf. Fig. 1, and are printable for effective physical-world attacks [35, 49]. On top of that, embedding the distorted optical flow into high-level recognition methods, e.g. for flow-based action recognition [9, 16], often corrupts the downstream application [17, 52].

^{*}Equal contribution.




Figure 1. Standard patch attack [35] (vanilla) and defense-aware attacks on FlowNetC’s [11] optical flow prediction. Left: Adversarial patch. Middle: Attacked image with applied defense (LGS or ILP, if any). Right: Optical flow. While both LGS and ILP defenses can defend against the vanilla patch attack [2, 35] (top) neither defense withstands defense-aware patch attacks (bottom).

To protect methods against the negative effects of adversarial patches, a straightforward defense concept is to first detect the patch and then render it harmless, e.g. by masking the former patch area [14, 30]. However, for classification, it was soon discovered that early defenses do not withstand attacks that take the defense mechanism into account [3, 10]. Such broken defenses are useless for practical applications because attackers aware of the method design (including potential defense mechanisms) can easily overcome them [3, 7, 46]. Among the broken defenses [10] for patch attacks on classification is Local Gradient Smoothing (LGS) [30], which also has been considered to defend optical-flow based action recognition pipelines [2]. Because LGS simply blackens the detected adversarial patch, Inpainting with Laplacian Prior (ILP) [2] was proposed. ILP

inpaints the patch region using neighborhood information to improve the classification accuracy for patch-attacked action recognition pipelines. However, the evaluation of LGS and ILP on the optical-flow component for action recognition in [2] has two major problems: First, it is unclear whether LGS or ILP withstand defense-aware attacks in the context of optical flow – given the results for LGS in classification [10], this is unlikely. And second, an evaluation of how these defenses affect the quality and robustness of optical flow methods is missing, which significantly impacts their practical applicability for all optical-flow-based problems. This work addresses both aspects by providing the first comprehensive analysis of detection defenses against patch attacks proposed in the context of optical flow.

Contributions. We make four contributions. (i) We develop *defense-aware patch attacks* on the ILP and LGS defense for optical flow estimation, by making the defenses differentiable (replacing gradient-free operations) and by avoiding patch detection by the defense (with tailored loss terms). (ii) Moreover, we *investigate the effectiveness* of ILP and LGS on a large set of optical flow methods. Surprisingly, the defenses not only lower the quality of benign (unattacked) predictions but also decrease the robustness for standard (vanilla) and defense-aware attacks – leaving no advantage of defended methods over undefended ones. (iii) Then, we find these *significant defense shortcomings to be caused* by the delocalized destruction of image information for benign scenes, which currently prevents viable detection defenses for optical flow estimation. (iv) Finally, we formulate evaluation advice for defenses on pixel-wise prediction tasks like optical flow, to help avoid common evaluation mistakes in future defense proposals.

2. Related work

Adversarial (patch) attacks on optical flow. Optical flow methods take a pair of input frames $I_1, I_2 \in \mathbb{R}^{M \times N \times 3}$ to predict the 2D vectors that describe the apparent motion, or optical flow $f \in \mathbb{R}^{M \times N \times 2}$ from I_1 to I_2 . Adversarial attacks then modify the input frames to *corrupt the optical flow* prediction. The first adversarial attacks on optical flow go back to Ranjan *et al.* [35] who considered patches [5]. Since then, patch attacks were extended to include transparencies [47], simultaneously harm depth estimators [50] or were used to attack flow-based action recognition [17]. Meanwhile, image-wide attacks on optical flow range from global [1, 39, 40] over semantically constrained attacks [20] to adversarial weather [37, 38]. Here, we investigate adversarial patches for being a threat in the physical world.

Adversarial defenses and their evaluation. Adversarial defense mechanisms are designed to protect methods against the perturbing effects of adversarial attacks. Typical defense strategies are adversarial training as a form

of data augmentation [13, 21, 36, 43], upstream strategies that filter perturbations from the inputs [14, 23, 30] and certified defenses that come with robustness guarantees [4, 10, 22, 33, 48]. However, many early defenses based on filtering operations were found to be ineffective if the attacker takes the defense mechanism into account [3, 7, 46]. Therefore, a defense’s effectiveness has to be shown under defense-aware adversarial attacks to justify its merit. In the process, one has to adequately (*i.e.* effectively) include the defense into the defense-aware attack: Prior work [3, 46] demonstrated in the context of *classification* that by neglecting this fact, many defenses appear unjustifiedly strong despite being evaluated with defense-aware attacks. Hence in this work, following the evaluation guidelines from [3, 46], we design the first defense-aware attacks on optical flow.

Defenses against adversarial patch attacks. Very few defenses are specialized to optical flow [2, 52]. Hence, we first discuss general adversarial patch defenses related to *classification*. Certifiable defenses against patch attacks on classification are provably robust [10, 22, 48], but often lead to smaller robustness improvements. Adversarial training [36] and architectural modification [29] have been also shown to improve the robustness against patch attacks. A last class of patch defenses aims to detect the patch in order to remove it [14, 23, 26, 30]. Among them, digital watermarking uses saliency maps [14], while Local Gradient Smoothing (LGS) detects anomalies in the input gradients [30]. Both use non-differentiable operations that hinder the backpropagation to train adversarial patches, but if these operations are replaced by differentiable ones [3], both defenses are ineffective against defense-aware attacks [10].

In the context of *optical flow*, LGS has been applied to defend optical-flow-based action recognition [2]. An optical-flow-specific improvement is Inpainting with Laplacian Prior (ILP) [2], which yields visually pleasing defended images. Also for action recognition, [52] proposed an optical-flow defense based on self-supervised counter-perturbations against noise-like perturbations. Since we focus on defenses against *patch attacks* for optical flow, this leaves LGS and ILP as potential methods. However, for such defenses, no analysis with defense-aware attacks has been performed, and neither have optical flow methods been considered independent of action recognition.

3. Defending optical flow with LGS and ILP

We begin by providing technical details for the LGS [30] and ILP [2] defenses. Both defenses detect the adversarial patch based on large image gradients and then replace these regions to remove the adversarial patch.

Patch detection. To detect the patch, the image is split into overlapping blocks $B = K \times K$ of size K and overlap O . Then, a subset of blocks containing potential adversar-




Figure 2. ILP (top) and LGS (bottom) defenses on unattacked (left) and attacked (right) images of the KITTI 2015 dataset [28]. Defenses degrade the visual quality, but LGS more than ILP.

ial modifications is selected. As adversarial patches often have large color changes (*cf.* Fig. 1), the gradient magnitude is accumulated for each block to identify blocks with the largest gradients. The gradient magnitude computation differs for ILP and LGS: While LGS considers first derivatives of the input image $I \in \mathbb{R}^{M \times N \times 3}$, ILP uses second derivatives, resulting in the gradient fields $G \in \mathbb{R}^{M \times N}$:

$$G_{\text{LGS}} = \|\nabla I\|, \quad (1)$$

$$G_{\text{ILP}} = \|\Delta I\|. \quad (2)$$

Normalizing gradients per image yields scale invariance:

$$\bar{G}_{i,j} = \frac{G_{i,j} - \min_{i,j \in M \times N} G_{i,j}}{\max_{i,j \in M \times N} G_{i,j} - \min_{i,j \in M \times N} G_{i,j}}. \quad (3)$$

Based on \bar{G} , adversarially modified pixels are marked: Per pixel (i, j) , we denote all enclosing blocks by $B_{(i,j)}$, and let these blocks vote whether the sum of their gradients exceeds a threshold $t \in [0, 1]$ (t is relative to the distribution of \bar{G}). If at least one block has large gradients, the respective pixel is marked as adversarial in a binary mask $M \in \mathbb{R}^{M \times N}$:

$$M_{i,j} = \begin{cases} 1 & \text{if } \exists B \in B_{(i,j)} : \sum_{k,l \in B} \bar{G}_{k,l} > t, \\ 0 & \text{else.} \end{cases} \quad (4)$$

Through this procedure, a block with large gradients causes all contained pixels to be marked as adversarial. To remove incorrectly marked (non-adversarial) pixels, ILP performs a reevaluation of candidates in M . After scaling with s_{ILP} , their gradients must exceed a threshold t_{ILP} to yield ILP’s final mask M_{ILP} , with \odot as pixel-wise multiplication:

$$M_{\text{ILP}} = M \odot \text{tr}(s_{\text{ILP}} \cdot \bar{G}_{\text{ILP}} > t_{\text{ILP}}). \quad (5)$$

Patch removal. Next, the defenses replace these potentially adversarial pixels from M for LGS and from M_{ILP} for ILP. LGS reduces the gradients in the detected area, which results in the modified image

$$I_{\text{LGS}} = (1 - \text{clip}_{[0,1]}(b_{\text{LGS}} \cdot \bar{G} \odot M)) \odot I, \quad (6)$$

where b_{LGS} is a smoothing parameter. If b_{LGS} is large, this darkens the selected adversarial pixels. ILP instead inpaints

the selected pixels with Telea’s algorithm [45] with radius r_{Telea} for more pleasing visual results. This smoothes colors from the edges of the selected areas into their center:

$$I_{\text{ILP}} = \text{Telea}(M_{\text{ILP}}, I, r_{\text{Telea}}), \quad (7)$$

Fig. 1 and Fig. 2 show LGS- and ILP-defended images. Our final hyperparameters for ILP and LGS are $K = 16$, $O = 8$, $t = 0.15$, $t_{\text{ILP}} = 0.5$, $s_{\text{ILP}} = 15$, $b_{\text{LGS}} = 15$ and $r_{\text{Telea}} = 5$; The supplement provides details on their selection.

4. Defense-aware patch attacks for optical flow

LGS and ILP defenses were only used to attack optical flow predictions for action recognition in a black-box way so far [2], meaning adversarial patches were trained without knowledge about the defense. According to best-practice for defense evaluation [3, 7, 46], the *defended* model must be evaluated under *defense-aware* attacks to show that it indeed offers protection. In the following, we develop white box patch attacks on the ILP and LGS defenses for optical flow. Our defense-aware attacks expand on Chiang *et al.* [10] who successfully attacked LGS for classification, but neither considered ILP nor the optical flow problem.

Gradient computations through the defense mechanism.

The defensive properties of LGS and ILP are based on *shattered gradients*, *i.e.* the use of mathematical operations with nonexistent gradients that prevent adversarial optimization [3, 31]. To still optimize adversarial patches in a defense-aware manner, the Backward Pass Differential Approximation (BPDA) [3] replaces these operations with differentiable approximations during backpropagation. The forward pass is executed normally. Within LGS and ILP, the problematic operations are the block-wise filtering steps (LGS and ILP), thresholding (ILP) and clipping operations (LGS), and the inpainting step (ILP). Below, we describe how they are approximated to enable backpropagation.

In the block-wise filtering step for LGS and ILP, *cf.* Eq. (4), the gradients do not exist for the conditional selection. To bypass them with BPDA [3], the filtering is replaced with the differentiable identity function, resulting in $\nabla M = 1$. The thresholding in ILP’s filtering has a similar problem, *cf.* Eq. (5), hence we also replace it with an identity function in the backward pass. For the clipping in LGS’s smoothing, *cf.* Eq. (6), the true gradient is one when the argument is in $[0, 1]$ and otherwise undefined. In practice, we find this operation responsible for most gradient shattering: Whenever the smoothing darkens values below zero, the clipping then sets them to zero, losing the gradient. Therefore, in the backpropagation, we approximate gradients with the identity if the value to clip is in $[0, 1]$ and with zero otherwise. As the ILP inpainting is very time-consuming, *cf.* Eq. (7), we treat it as being gradient-free. Similar to the clipping approximation, we bypass it

with an identity operation for non-inpainted pixels and a zero-gradient for inpainted ones. To overcome optimization problems for zero-gradients in the clipping- and inpainting approximations, we introduce additional loss terms to improve the patch in areas with no gradient information.

Defense-aware loss functions. Optimizing defense-aware adversarial patches requires a loss function that encourages patches with a perturbing effect on the optical flow output. As baseline loss that defines the overall goal for the patch attack, we use the Average Cosine Similarity (ACS) which was used to train adversarial patches on optical flow in [35]. It encourages adversarial patches to invert the original optical flow prediction f to yield the adversarial flow \check{f} :

$$\mathcal{L}^{\text{ACS}}(f, \check{f}) = \frac{1}{NM} \sum_{i,j \in M \times N} \frac{\langle f_{i,j}, \check{f}_{i,j} \rangle}{\|f_{i,j}\|_2 \|\check{f}_{i,j}\|_2}. \quad (8)$$

Besides the ACS, another loss term is required to overcome the zero-gradients of BPDA. While the differentiable LGS and ILP approximations allow optimizing adversarial patches, these patches may still be detected by the defenses and hence be stopped from perturbing the flow. Therefore, we use loss terms to penalize large gradient magnitudes in the patches. To optimize defense-aware patches P , we therefore penalize first-order gradients $\|\nabla P\|$ for LGS and second-order gradients $\|\Delta P\|$ for ILP-awareness:

$$\mathcal{L}^{\text{LGS}}(f, \check{f}, P) = \mathcal{L}^{\text{ACS}}(f, \check{f}) + \alpha \|\nabla P\|, \quad (9)$$

$$\mathcal{L}^{\text{ILP}}(f, \check{f}, P) = \mathcal{L}^{\text{ACS}}(f, \check{f}) + \alpha \|\Delta P\|. \quad (10)$$

The parameter α balances the loss terms and is set to $\alpha = 1e-8$. With small gradient magnitudes, the patches are likely below the filtering threshold as it is relative to the remaining image gradients. This way, they evade the defenses and affect the optical flow output as in Fig. 1.

In the ACS implementation, we exclude the patch area from the computation. This measures to which extent the patch modifies the optical flow outside its direct area, *i.e.* it assesses the de-localized impact per patch. This is because one may take two points of view on the role of the patch: In the first view, the patch is an image part, with a zero ground-truth flow at the patch area. In the second view, the patch is an attack part, and defenses should mitigate its effect and restore the ground truth flow in its area. To refrain from assuming a ‘‘correct’’ optical flow for the patch, we exclude the patch region from our loss.

Defense-aware patch optimization. We test two different methods for optimization: The Iterative Fast Gradient Sign Method (I-FGSM) [21] and Stochastic Gradient Descent (SGD). To ensure a valid color range of the patch P after optimization, we consider clipping the values to their valid range in $[0, 1]$ after each update [35] and a change of variables (CoV) via tanh to optimize the values in $[-\infty, \infty]$ before transforming them back into the valid range [8, 39].

5. Metrics for defended quality and robustness

Including a defense to protect an existing method against attacks effectively creates a new method that consists of the original method plus defense D . Hence, we have to evaluate the quality and robustness of this new method instead of the original defense-free approach’s metrics [7, 46].

Quality. To evaluate the quality of defended optical flow methods, one typically measures the average endpoint error (EPE) between the ground-truth flow f^* and the predicted flow f , where low errors indicate high quality:

$$\text{EPE}(f^*, f) = \frac{1}{MN} \sum_{i \in M \times N} \|f_i^* - f_i\|_2. \quad (11)$$

Robustness. To evaluate robustness, we use the methodology from [39] and measure the distance between the benign and the attacked flow prediction. This quantifies how much an attack changes a method’s output and is motivated by the Lipschitz continuity of functions. Due to the previously discussed two views on the adversarial patch, we evaluate the EPE for all pixels *outside* P , quantifying the patch’s negative effect outside its immediate area. We denote the benign flow prediction of a method defended with D as f_D and its prediction under attack by A as f_D^A . Then, the robustness is

$$\text{EPE}_P(f_D, f_D^A) = \frac{1}{MN - P} \sum_{i \in M \times N \setminus P} \|(f_D)_i - (f_D^A)_i\|_2 \quad (12)$$

with low values for robust methods, as attacked predictions outside the patch should coincide with the unattacked ones.

Quality and robustness for pipelines. When a method is defended with a defense D and attacked with an attack that is defense-aware towards D , we call the setup a *full pipeline* with defense D . Its quality is $Q_D = \text{EPE}(f^*, f_D)$ and the resulting pipeline robustness is $R_D^D = \text{EPE}_P(f_D, f_D^D)$.

6. Experiments

We now assess how defenses against patch attacks impact the quality and robustness of optical flow methods. We begin by evaluating the quality of defended methods and then separately assess their robustness against patch attacks. Afterward, we jointly analyze both, quality and robustness, to find them being negatively impacted by defenses against patch attacks. Finally, we explore the reasons for their poor performance. All attacks and defenses are implemented with PyTorch [32], and available at <https://github.com/cv-stuttgart/DetectionDefenses>.

Optical flow methods and attack setups. As optical flow methods, we select FlowNetC (FNC) [11], SpyNet [34] and PWCNet (PWC) [42] as milestone architectures in flow estimation, RAFT [44], GMA [18] and FlowFormer




Figure 3. Overview of adversarial patches with size 100 for vanilla, ILP- and LGS-aware patch attacks against all tested networks.

(FF) [15] as current state-of-the-art, and FlowNetCRobust (FNCR) [40] as it improves FlowNetC’s patch robustness.

We generate effective adversarial patches by optimizing learning rates, box constraints and optimizer choice for each optical flow network. As optimizers, we consider I-FGSM [21] and SGD, learning rates from 0.1-100 (optimizer-dependent), box constraints via change of variables (CoV) or clipping. Following the protocol for adversarial patches from [35], we optimize patches of size 100 using KITTI Raw [12] and evaluate on KITTI train [28]. Evaluations using Sintel [6], Driving [25], HD1K [19] and Spring [27] are shown in the supplement. For defense-aware patches, we choose $\alpha = 1e-8$ for the loss function in Eq. (9) and Eq. (10). Per flow method, we generate vanilla adversarial patches (without defense awareness, as in [35]) and defense-aware patches for LGS [30] or ILP [2]. We train 4 patches per parameter combination, and average the evaluated metrics. Among the parameters, we select the strongest adversarial configuration for the worst robustness.

The full evaluation of the best parameters per flow method and defense strategy is in the supplement. Fig. 3 visualizes the most effective patches. Both defenses detect high gradient magnitudes, causing smooth defense-aware patches with small derivatives. Interestingly, defense-aware patches for methods like RAFT, GMA or FlowFormer contain high-frequency noise. This calls for detection rather than evasion by ILP or LGS, which we explore in Sec. 6.4.

6.1. Quality of defended optical flow methods

To begin our investigation of defenses D, we assess the quality of defended and undefended optical flow methods on unattacked input frames. Tab. 1 lists the endpoint errors $EPE(f^*, f_D)$ on KITTI train, where the ground truth flow is available. Across all optical flow methods, we find the lowest errors when no defense is applied; Both ILP and LGS

Table 1. Quality $Q_D = EPE(f^*, f_D)$ for optical flow pipelines with defense D on the KITTI train dataset [28]; Best quality is **bold**. All defenses lead to a worse quality on unattacked frames.

Defense		FNCR	FNCR	SpyNet	PWC	RAFT	GMA	FF
None	Q	15.42	11.10	24.03	13.26	0.63	0.61	0.62
LGS	Q_{LGS}	16.70	13.13	25.15	14.61	1.42	1.55	1.42
ILP	Q_{ILP}	16.46	12.77	24.74	14.52	1.36	1.39	1.30

lead to larger errors. But for accurate methods, the errors rise more than for less accurate ones, *i.e.* by 156% for GMA and 4% for SpyNet, using LGS vs. no defense. On average, ILP increases the error less than LGS, *i.e.* by 129% instead of 156% for GMA, compared to no defense. ILP performs better due to its more sophisticated image restoration, which adds pixel-wise filtering with inpainting rather than smoothing, *cf.* Fig. 2. In the figure, applying ILP and LGS to unattacked images visually degrades them, leading to worse predicted flows. Overall, detect-and-remove defenses lower the accuracy on unattacked frames, as they strongly affect the image quality, which harms the flow quality.

6.2. Robustness under defense-aware patch attacks

To study the robustness of defended flow methods under defense-aware attacks, we measure $R_D^A = EPE_P(f_D, f_D^A)$ for all combinations of defenses D (None, LGS and ILP) and attacks A (vanilla, LGS- and ILP-aware). Tab. 2 gives the full results. In the analysis process, we (i) evaluate whether defense-aware attacks bypass the defenses, and then (ii) identify the most effective defense for each attack.

Most effective attack per defense. First, we evaluate if our defense-aware attacks evade the defenses. In practice,

Table 2. Robustness scores for all combinations of defended methods and defense-aware attacks on optical flow methods on KITTI train [28]. For a given defense D and attack A, the robustness is defined as $R_D^A = EPE_P(f_D, f_D^A)$. Per attack, the robustness values of the best defense are **bold**. Per defense, the robustness values for the attack it is most vulnerable to are underlined. Full pipelines are highlighted in gray, and provide the corresponding robustness values to the quality scores from Tab. 1.

Attack type	Defense	FNC	FNCR	SpyNet	PWC	RAFT	GMA	FF
Vanilla	None R^{Van}	73.74	1.78	1.48	2.17	0.33	0.56	0.57
	LGS $R_{\text{LGS}}^{\text{Van}}$	3.75	2.97	3.97	3.34	1.45	1.31	1.30
	ILP $R_{\text{ILP}}^{\text{Van}}$	4.66	3.11	3.34	3.29	1.43	0.99	1.41
+LGS (LGS-aware)	None R^{LGS}	50.46	0.46	1.40	2.10	0.25	0.27	0.44
	LGS $R_{\text{LGS}}^{\text{LGS}}$	<u>23.36</u>	<u>3.27</u>	4.05	4.13	1.46	<u>1.60</u>	1.67
	ILP $R_{\text{ILP}}^{\text{LGS}}$	23.04	<u>4.21</u>	3.32	3.35	1.48	<u>1.54</u>	1.80
+ILP (ILP-aware)	None R^{ILP}	56.56	1.02	1.45	2.16	0.20	0.26	0.45
	LGS $R_{\text{LGS}}^{\text{ILP}}$	10.99	3.68	3.03	4.06	1.47	1.57	1.68
	ILP $R_{\text{ILP}}^{\text{ILP}}$	<u>55.26</u>	3.25	<u>3.36</u>	<u>4.25</u>	<u>1.49</u>	1.51	<u>1.81</u>

this corresponds to choosing a defense to observe how the defended model fares against different attacks. For a fixed defense D in Tab. 2 we underline the worst robustness, *i.e.* the most effective attack. Hence we compare R_D^{Van} , R_D^{LGS} and R_D^{ILP} (*e.g.* the 2nd line in each block for D=LGS).

Without defenses every method is most vulnerable towards the vanilla attack: $R^{\text{Van}} \geq R^{\text{LGS}}, R^{\text{ILP}}$. This is plausible, as LGS- and ILP-aware attacks impose additional constraints on the patches, which impairs their effectiveness for an undefended model. For defended models, the corresponding defense-aware attacks are often most effective, *e.g.* $R_{\text{LGS}}^{\text{LGS}}$ is largest for the LGS-defended FlowNetC, FlowNetCRobust, SpyNet and GMA. This confirms that our adaptive attacks are truly defense-aware, as they are most effective on the defended models, *i.e.* $R_D^D \geq R_A^D$ for the majority of models. Still, in some cases, an ILP-aware attack performs better on an LGS-defended model and vice versa. This indicates transferable patches for LGS and ILP, as the differences are small in these cases, *e.g.* LGS-defended RAFT scores $R_{\text{LGS}}^{\text{LGS}} = 1.46$ and $R_{\text{ILP}}^{\text{LGS}} = 1.47$.

Overall, our defense-aware attacks are most effective w.r.t. the respective defended models, which validates their design and implementation. Likewise, the defense-aware attacks are less effective on other defenses, as inappropriate constraints hinder the patch’s effectiveness. Nonetheless, we find that LGS- and ILP-aware patches are transferable.

Most effective defense per attack. Next, we analyze the most effective defense for a given attack; or in other words, which defense withstands most attacks. Per fixed attack A in Tab. 2, we **boldface** the best robustness per network, comparing R^A , R_{LGS}^A and R_{ILP}^A with differing defenses.

Focusing first on the vanilla attack (Tab. 2 block 1), the undefended robustness R^{Van} strongly differs. FlowNetC is particularly vulnerable [2, 35, 39, 40] due to a limited field of vision that was expanded in FlowNetCRobust [40] and improved its robustness by 97%, making it comparable to




Figure 4. Optical flow estimations for selected methods on a KITTI frame that is unattacked (left) and attacked with the vanilla patch attack (right). Blue circles indicate the patch location. An overview of all optical flow methods is in the supplement.

SpyNet or PWCNet. Most robust against vanilla patch attacks are the state-of-the-art methods RAFT, GMA and FlowFormer. Their robustness appears linked to their quality, as they detect the static patches in Fig. 4, correctly estimating the zero motion. This retains correct flow predictions around the patch and results in low robustness scores.

For methods that are robust against vanilla attacks without defense, *i.e.* all except FlowNetC, defending harms their robustness scores: $R^A < R_{\text{LGS}}^A, R_{\text{ILP}}^A$, independent of the attack A (vanilla, LGS- or ILP-aware). This renders the defenses ineffective, as improving robustness against attacks is their sole purpose. In contrast, defending FlowNetC improves its robustness against vanilla attacks from 73.74 to 3.75 with LGS and 4.66 with ILP. For action recognition, a similar improvement was seen with FlowNetC [2], but compared to other methods, FlowNetC is not robust even when defended and therefore should not be used.

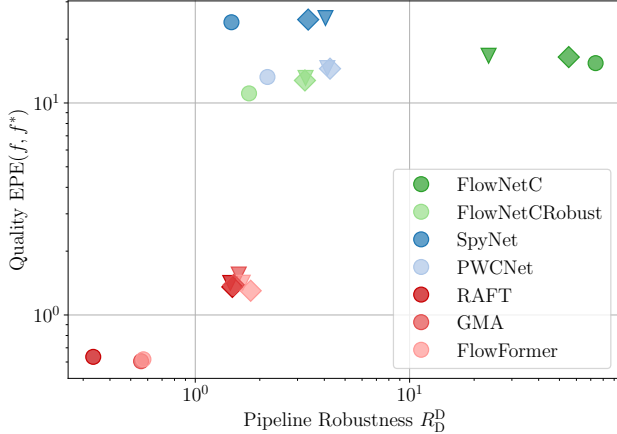


Figure 5. Quality vs. robustness of flow networks on KITTI train in a double logarithmic plot. An ideal method would be in the origin. Undefended networks are circles \circ , networks defended with LGS are triangles ∇ and networks defended with ILP are diamonds \diamond . ILP and LGS deteriorate quality and robustness.

All in all, the reported robustness enhancement through ILP and LGS in [2] can not be confirmed for our large test body of optical flow methods. Instead, we find that LGS and ILP defenses harm the robustness of competitive optical flow methods for all tested patch attacks.

6.3. Quality and robustness for defended methods

After separately considering quality and robustness of defended methods under adversarial patch attacks, we now jointly analyze both aspects. Perfect defenses decrease the vulnerability to adversarial attacks without negatively impacting the quality. As models, we consider all tested flow methods with no defense, LGS or ILP. Their quality is taken from Tab. 1. Their corresponding pipeline robustness, *i.e.* defended model’s robustness under the respective defense-aware attacks, which is highlighted in gray in Tab. 2.

For all optical flow methods, Fig. 5 visualizes the quality-robustness pairs per defense, *e.g.* Q_{LGS} with R_{LGS}^{LGS} . An ideal method with low scores for quality and robustness would be positioned at the origin. An improvement in robustness moves the defended point to the left, ideally without decreasing quality. For all methods except FlowNetC, the undefended standard model (\circ) is closest to the origin and therefore offers the best robustness *and* the best quality, without any trade-off. Using LGS (∇) or ILP (\diamond) defenses worsen both metrics to a similar extent. The only outlier is FlowNetC, where both defenses improve the robustness while keeping the quality nearly constant, with larger improvements for LGS than for ILP. Overall, our investigation shows that almost all optical flow methods are harmed by the detect-and-remove defenses ILP and LGS, as they worsen method quality and robustness alike.



Figure 6. Effect of the LGS defenses on KITTI [28] frames (left) and the resulting optical flow prediction with RAFT [44] (right). Black areas in the input frame are filtered by the LGS defense. Blue circles mark the area of the adversarial patch, the red boxes highlight an area with prominent differences in the flow predictions. Note that the robustness calculation omits the blue circle.

6.4. Flaws explained: Manual patch attack

From Tab. 2 we saw significant robustness reductions for high-quality methods like RAFT, GMA or FlowFormer when defended with ILP or LGS. Yet, the reductions are caused by high-frequency defense-aware patches, *cf.* Fig. 3, which seems to contradict the optimization for smoothness to evade detection by the ILP and LGS gradient filtering.

To understand this behavior, we compare the flows entering into the robustness calculation – the unattacked flow f_D of the defended method and the flow f_D^A after applying a defense-aware attack to the defended model. Fig. 6 shows RAFT’s original prediction (unattacked, no defense, Row 1) together with flows for the LGS-defended version. Comparing defended and undefended flows, *e.g.* the car in the red box, the flow from the *unattacked* LGS-defended RAFT (f_{LGS} , Row 2) is very erroneous compared to the *attacked* LGS-defended flow (f_{LGS}^{LGS} , Row 3). In other words, the gradient filtering of the defense destroys important visual information throughout the image, which yields low-quality optical flow predictions in the absence of adversarial attacks. If the alterations in an unattacked image are scattered throughout its domain, a patch attack can maximize flow changes by aggregating alterations in a single location, *i.e.* the patch itself. Incidentally, this *improves* the optical flow prediction in large areas (Rows 1, 3; Red box).

Therefore, we hypothesize that the bad robustness scores of high-quality methods are driven by large distortions in unattacked frames caused by the defense. To test this, we design a manual patch (see Fig. 7) consisting of a checkerboard pattern to maximize first- and second derivatives. With the manual patch, we then attack defended and undefended optical flow methods. Their robustness in Tab. 3 confirms our hypothesis. For undefended methods, the




Figure 7. Visual comparison of patches obtained for vanilla, LGS-aware and manual patch attack on FlowFormer [15]. The manual patch imitates high derivatives in the LGS-aware adversarial patch.

Table 3. Robustness $R_D^{\text{Man}} = \text{EPE}_P(f_D, f_D^{\text{Man}})$ against a manual patch attack (Man) of optical flow methods with different defenses D. The best robustness is **bold**.

Defense	FNC	FNCR	SpyNet	PWC	RAFT	GMA	FF
None R_D^{Man}	1.19	0.51	1.15	0.90	0.19	0.25	0.44
LGS $R_{\text{LGS}}^{\text{Man}}$	3.69	3.26	3.86	3.40	1.45	1.56	1.61
ILP $R_{\text{ILP}}^{\text{Man}}$	3.84	3.35	3.18	3.49	1.53	1.57	1.86

patch hardly affects the optical flow prediction. For defended methods, however, the robustness significantly degrades, even though the patch impacts the defense, not the network. Also, we obtain similar results to those of defense-aware attacks in Tab. 2, where robustness degrades most significantly for accurate methods, *e.g.* RAFT, GMA and FlowFormer. However, because FlowNetC’s robustness $\text{EPE}_P(f_D, f_D^{\text{A}})$ is driven by the attacked flow f_D^{A} rather than the unattacked f_D , *cf.* Fig. 4, defenses can succeed as every attacked-flow improvement directly serves the robustness.

In summary, the sub-par quality and robustness of defended high-quality methods are a direct consequence of the defense itself, which causes visual distortions in unattacked frames. These distortions not only reduce the quality of benign frames but also are the empirically-confirmed cause for the low robustness scores of high-quality methods.

7. Discussion

We take a moment to condense the findings from analyzing defended optical flow methods into actionable evaluation advice, to encourage meaningful defense evaluations in the future. While evaluation advice has been formulated before and should be adhered [3, 7, 24, 46], we refresh some points, reinforce their importance and add discussions specific to pixel-wise prediction tasks like optical flow.

Quality changes with defense. Defending a method creates a modified method and thus modifies its quality characteristics. Therefore, the quality of the defended method Q_D should be explicitly reported. Particularly for pixel-wise prediction tasks, subtle changes in the inputs can cause

significant output changes over large areas, making it indispensable to *reevaluate the defended quality*.

Use defense-aware attacks. Every defense proposal must be evaluated with a sufficiently *strong adaptive attack* [3, 7, 46] and report the pipeline robustness R_D^{D} . Showing that it withstands adversarial samples for the original method is *not enough*. While many defense-circumvention strategies for classification [46] may apply, individual tuning to pixel-wise prediction is needed for strong adaptive attacks.

Components matter. Defenses for specific components of complex methods should be evaluated on the *specific part*, not only on the full method. *E.g.* when defending optical flow for action recognition [2, 52], defended quality Q_D and defense-aware robustness R_D^{D} should be reported for the flow component. Otherwise, the defense effectiveness and method sensitivity towards the component are entangled.

8. Limitations

This work solely focuses on detect-and-remove defenses for optical flow estimation. Hence, it covers neither defenses for problems unrelated to optical flow, nor optical flow defenses against non-patch attacks (of which none were published so far). Our findings that detection defenses do not protect against patch attacks could transfer to future defenses based on the gradient magnitude, as we found ILP-aware patches to transfer to LGS-defended methods and vice versa. Nonetheless, defending optical flow may be possible with more specialized techniques.

9. Conclusion

We investigated detect-and-remove defenses against adversarial patch attacks on optical flow methods. To this end, we designed defense-aware patches that avoid detection by LGS and ILP defenses, allowing us to break both defenses on a large variety of optical flow methods. On top of that, we found that both defenses reduce the optical flow quality and even failed to increase the robustness against standard (*i.e.*, not defense-aware) attacks. We could attribute this discouraging performance to the severe image quality degradation resulting from pixel replacements in the defenses. As image quality is crucial for pixel-wise motion estimation, this illustrates that defenses for classification methods, like LGS, do not automatically protect optical flow. Consequently, flow pipelines’ robustness and quality must be thoroughly investigated for every defense, to promote trust instead of making empty promises.

Acknowledgments. We thank Filip Ilic for helpful discussions. The International Max Planck Research School for Intelligent Systems supports JS. Funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – Project-ID 251654672 – TRR 161 (B04).

References

- [1] Shashank Agnihotri, Steffen Jung, and Margret Keuper. CosPGD: A unified white-box adversarial attack for pixel-wise prediction tasks. *arXiv:2302.02213*, 2023. [2](#)
- [2] Adithya Prem Anand, H. Gokul, Harish Srinivasan, Pranav Vijay, and Vineeth Vijayaraghavan. Adversarial patch defense for optical flow networks in video action recognition. In *Proc. IEEE International Conference on Machine Learning and Applications (ICMLA)*, pages 1289–1296, 2020. [1](#), [2](#), [3](#), [5](#), [6](#), [7](#), [8](#), [12](#)
- [3] Anish Athalye, Nicholas Carlini, and David Wagner. Obfuscated gradients give a false sense of security: Circumventing defenses to adversarial examples. In *Proc. International Conference on Learning Representations (ICML)*, pages 274–283, 2018. [1](#), [2](#), [3](#), [8](#)
- [4] Wieland Brendel and Matthias Bethge. Approximating CNNs with bag-of-local-features models works surprisingly well on ImageNet. In *Proc. International Conference on Learning Representations (ICML)*, 2019. [2](#)
- [5] Tom B. Brown, Dandelion Mané, Aurko Roy, Martín Abadi, and Justin Gilmer. Adversarial patch. In *NeurIPS Workshop on Machine Learning and Computer Security (NeurIPS-MLCS)*, 2017. [2](#)
- [6] D. J. Butler, J. Wulff, G. B. Stanley, and M. J. Black. A naturalistic open source movie for optical flow evaluation. In *Proc. European Conference on Computer Vision (ECCV)*, pages 611–625. Springer, 2012. [5](#), [18](#), [19](#), [22](#)
- [7] Nicholas Carlini and David Wagner. Adversarial examples are not easily detected: Bypassing ten detection methods. In *ACM Workshop on Artificial Intelligence and Security (AiSec)*, pages 3–14, 2017. [1](#), [2](#), [3](#), [4](#), [8](#)
- [8] Nicholas Carlini and David Wagner. Towards evaluating the robustness of neural networks. In *IEEE Symposium on Security and Privacy (SP)*, pages 39–57, 2017. [4](#)
- [9] Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? a new model and the Kinetics dataset. In *Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. [1](#)
- [10] Ping-Yeh Chiang, Renkun Ni, Ahmed Abdelkader, Chen Zhu, Christoph Studor, and Tom Goldstein. Certified defenses for adversarial patches. In *Proc. International Conference on Learning Representations (ICLR)*, 2020. [1](#), [2](#), [3](#)
- [11] Alexey Dosovitskiy, Philipp Fischer, Eddy Ilg, Philip Hausser, Caner Hazirbas, Vladimir Golkov, Patrick van der Smagt, Daniel Cremers, and Thomas Brox. FlowNet: Learning optical flow with convolutional networks. In *Proc. IEEE/CVF International Conference on Computer Vision (ICCV)*, 2015. [1](#), [4](#), [5](#), [12](#), [14](#)
- [12] Andreas Geiger, Philip Lenz, Christoph Stiller, and Raquel Urtasun. Vision meets robotics: The KITTI dataset. *International Journal in Robotics Research (IJRR)*, 32(11):1231–1237, 2013. [5](#), [19](#)
- [13] Thomas Gittings, Steve Schneider, and John Collomosse. Vax-a-Net: Training-time defence against adversarial patch attacks. In *Proc. Asian Conference on Computer Vision (ACCV)*, 2020. [2](#)
- [14] Jamie Hayes. On visible adversarial perturbations & digital watermarking. In *Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 1597–1604, 2018. [1](#), [2](#)
- [15] 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. [5](#), [8](#), [12](#), [13](#), [14](#)
- [16] Filip Ilic, Thomas Pock, and Richard P. Wildes. Is appearance free action recognition possible? In *Proc. European Conference on Computer Vision (ECCV)*, pages 156–173, 2022. [1](#)
- [17] Nathan Inkawhich, Matthew Inkawhich, Yiran Chen, and Hai Li. Adversarial attacks for optical flow-based action recognition classifiers. *arXiv:1811.11875*, 2018. [1](#), [2](#)
- [18] 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. [4](#), [5](#), [12](#), [14](#)
- [19] Daniel Kondermann, Rahul Nair, Katrin Honauer, Karsten Krispin, Jonas Andrulis, Alexander Brock, Burkhard Gussefeld, Mohsen Rahimimoghaddam, Sabine Hofmann, Claus Brenner, and Bernd Jähne. The HCI benchmark suite: Stereo and flow ground truth with uncertainties for urban autonomous driving. In *Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 19–28, 2016. [5](#), [18](#), [19](#), [21](#)
- [20] Tom Koren, Lior Talker, Michael Dinerstein, and Ran Vitek. Consistent semantic attacks on optical flow. In *Proc. Asian Conference on Computer Vision (ACCV)*, pages 1658–1674, 2022. [2](#)
- [21] Alexey Kurakin, Ian J. Goodfellow, and Samy Bengio. Adversarial machine learning at scale. In *Proc. International Conference on Learning Representations (ICLR)*, 2017. [2](#), [4](#), [5](#), [12](#)
- [22] Alexander Levine and Soheil Feizi. (De)randomized smoothing for certifiable defense against patch attacks. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin, editors, *Proc. Conference on Neural Information Processing Systems (NeurIPS)*, pages 6465–6475, 2020. [2](#)
- [23] Jiang Liu, Alexander Levine, Chun Pong Lau, Rama Chellappa, and Soheil Feizi. Segment and complete: Defending object detectors against adversarial patch attacks with robust patch detection. In *Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 14973–14982, 2022. [2](#)
- [24] Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adria Vladu. Towards deep learning models resistant to adversarial attacks. In *Proc. International Conference on Learning Representations (ICML)*, pages 1–10, 2018. [8](#)
- [25] Nikolaus Mayer, Eddy Ilg, Philip Hausser, Philipp Fischer, Daniel Cremers, Alexey Dosovitskiy, and Thomas Brox. A large dataset to train convolutional networks for disparity, optical flow, and scene flow estimation. In *Proc. IEEE/CVF*

- Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4040–4048, 2016. [5](#), [18](#), [19](#), [22](#)
- [26] Michael McCoyd, Won Park, Steven Chen, Neil Shah, Ryan Roggenkemper, Minjune Hwang, Jason Xinyu Liu, and David Wagner. Minority reports defense: Defending against adversarial patches. In *Proc. International Conference on Applied Cryptography and Network Security Workshops (ACNSW)*, pages 564–582, 2020. [2](#)
- [27] Lukas Mehl, Jenny Schmalzfuss, Azin Jahedi, Yaroslava Naliwayko, and Andrés Bruhn. Spring: A high-resolution high-detail dataset and benchmark for scene flow, optical flow and stereo. In *Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4981–4991, 2023. [5](#), [18](#), [19](#), [21](#)
- [28] Moritz Menze and Andreas Geiger. Object scene flow for autonomous vehicles. In *Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3061–3070, 2015. [3](#), [5](#), [6](#), [7](#), [18](#), [19](#)
- [29] Norman Mu and David Wagner. Defending against adversarial patches with robust self-attention. In *ICML Workshop on Uncertainty and Robustness in Deep Learning (ICML-UDL)*, 2021. [2](#)
- [30] Muzammal Naseer, Salman Khan, and Fatih Porikli. Local gradients smoothing: Defense against localized adversarial attacks. In *Proc. IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pages 1300–1307, 2019. [1](#), [2](#), [5](#), [12](#)
- [31] Nicolas Papernot, Patrick McDaniel, Ian Goodfellow, Somesh Jha, Z. Berkay Celik, and Ananthram Swami. Practical black-box attacks against machine learning. In *ACM Asia Conference on Computer and Communications Security (ASIA-CCS)*, pages 506–519, 2017. [3](#)
- [32] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. PyTorch: An imperative style, high-performance deep learning library. In *Proc. Conference on Neural Information Processing Systems (NeurIPS)*, pages 8024–8035, 2019. [4](#)
- [33] Aditi Raghunathan, Jacob Steinhardt, and Percy Liang. Certified defenses against adversarial examples. In *Proc. International Conference on Learning Representations (ICLR)*, 2018. [2](#)
- [34] 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)*, 2017. [4](#), [5](#), [12](#), [14](#)
- [35] Anurag Ranjan, Joel Janai, Andreas Geiger, and Michael J. Black. Attacking optical flow. In *Proc. IEEE/CVF International Conference on Computer Vision (ICCV)*, 2019. [1](#), [2](#), [4](#), [5](#), [6](#), [12](#), [19](#)
- [36] Sukrut Rao, David Stutz, and Bernt Schiele. Adversarial training against location-optimized adversarial patches. In *Proc. IEEE/CVF International Conference on Computer Vision Workshops (ECCVW)*, pages 429–448, 2020. [2](#)
- [37] Jenny Schmalzfuss, Lukas Mehl, and Andrés Bruhn. Attacking motion estimation with adversarial snow. *ECCV Workshop on Adversarial Robustness in the Real World (ECCV-AROW)*, 2022. [2](#)
- [38] Jenny Schmalzfuss, Lukas Mehl, and Andrés Bruhn. Distracting downpour: Adversarial weather attacks for motion estimation. In *Proc. IEEE/CVF International Conference on Computer Vision (ICCV)*, 2023. [2](#)
- [39] Jenny Schmalzfuss, 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. [2](#), [4](#), [6](#), [13](#)
- [40] Simon Schrodi, Tonmoy Saikia, and Thomas Brox. Towards understanding adversarial robustness of optical flow networks. In *Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8916–8924, 2022. [2](#), [5](#), [6](#), [12](#), [14](#)
- [41] Deqing Sun, Charles Herrmann, Fitsum Reda, Michael Rubinstein, David J. Fleet, and William T. Freeman. Disentangling architecture and training for optical flow. In *Proc. European Conference on Computer Vision (ECCV)*, pages 165–182, 2022. [18](#)
- [42] Deqing Sun, Xiaodong Yang, Ming-Yu Liu, and Jan Kautz. PWC-Net: CNNs for optical flow using pyramid, warping, and cost volume. In *Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. [4](#), [5](#), [12](#), [14](#)
- [43] Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. Intriguing properties of neural networks. In *Proc. International Conference on Learning Representations (ICLR)*, 2014. [2](#)
- [44] 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. [4](#), [5](#), [7](#), [12](#), [14](#)
- [45] Alexandru Telea. An image inpainting technique based on the fast marching method. *Journal of Graphics Tools (JGT)*, 9(1):23–34, 2004. [3](#)
- [46] Florian Tramer, Nicholas Carlini, Wieland Brendel, and Aleksander Madry. On adaptive attacks to adversarial example defenses. In *Proc. Conference on Neural Information Processing Systems (NeurIPS)*, pages 1633–1645, 2020. [1](#), [2](#), [3](#), [4](#), [8](#)
- [47] Benjamin Wortman. Hidden patch attacks for optical flow. In *ICML Workshop on Adversarial Machine Learning (ICML-AdvML)*, 2021. [1](#), [2](#)
- [48] Chong Xiang, Arjun Nitin Bhagoji, Vikash Sehwal, and Prateek Mittal. PatchGuard: A provably robust defense against adversarial patches via small receptive fields and masking. In *Proc. USENIX Security Symposium*, 2021. [2](#)
- [49] Koichiro Yamanaka, Ryutaroh Matsumoto, Keita Takahashi, and Toshiaki Fujii. Adversarial patch attacks on monocular depth estimation networks. *IEEE Access*, 8:179094–179104, 2020. [1](#)
- [50] Koichiro Yamanaka, Keita Takahashi, Toshiaki Fujii, and Ryutaroh Matsumoto. Simultaneous attack on CNN-based monocular depth estimation and optical flow estimation. *IE-*

ICE Transactions on Information and Systems, (5):785–788, 2021. [1](#), [2](#)

- [51] Gengshan Yang and Deva Ramanan. Volumetric correspondence networks for optical flow. In *Proc. Conference on Neural Information Processing Systems (NeurIPS)*, volume 32. Curran Associates, Inc., 2019. [18](#), [19](#), [22](#)
- [52] Lingyu Zhang, Chengzhi Mao, Junfeng Yang, and Carl Vondrick. Adversarially robust video perception by seeing motion. *arXiv:2212.07815*, 2022. [1](#), [2](#), [8](#)

A. Supplementary Material

In our evaluations, we consider the optical flow methods FlowNetC (FNC) [11], FlowNetCRobust (FNCR) [40], PWCNet (PWC) [42], SpyNet¹ [34], RAFT [44], GMA [18] and FlowFormer (FF) [15].

A.1. Defense hyperparameter evaluation

For the LGS [30] and ILP [2] defenses, we identify those hyperparameters that lead to the most effective defense against the vanilla patch attack [35] on FlowNetC [11]. The hyperparameters under consideration are the block size K , block overlap O and the block filtering threshold t , which are used in ILP and LGS. For LGS, we further consider the smoothing parameter b_{LGS} . For ILP, we consider the scaling s_{ILP} , inpainting radius r_{Talea} and the threshold t_{ILP} . Out of those, we directly set $r_{Talea} = 5$ and $t_{ILP} = 0.5$, which are the values from Anand *et al.* [2] that also produced good results for our experiments. For the other parameters, we perform a parameter study that jointly evaluates the parameter pairs K vs. O (for LGS and ILP), t vs. b_{LGS} (for LGS) and t vs. s_{ILP} (for ILP).

Per parameter combination, we evaluate the robustness of the defended FlowNetC against the vanilla attack via $EPE(f, f_D^{Van})$ (Fig. A1 left, small values indicate good robustness) and also quantify how much the defense changes the flow prediction for unattacked frames via $EPE(f, f_D)$ (Fig. A1 right, small values indicate that defense does not change the flow prediction on benign samples). Fig. A1 shows the plots for both metrics and all parameter pairs on LGS and ILP. To select the parameters, for each parameter pair we pick values that lead to small values in both metrics (dark colors in plots for $EPE(f, f_D^{Van})$ and $EPE(f, f_D)$), because then the defense protects against vanilla attacks but at the same time does not change the flow prediction on unattacked samples. Thus, we select the parameters $K = 16$, $O = 8$, $s_{ILP} = 15$, $b_{LGS} = 15$ and $t = 0.15$, which offer the best trade-off between the two metrics. Note that for the K vs. O plots (Fig. A1, Row 1 and 3), the dark area with low values in the upper left corner is unfeasible, because the block overlaps O can not be larger than the blocksize K . Overall, our optimized parameters differ slightly from the literature values: For LGS, the original publication [30] used $K = 15$, $O = 5$ and $b_{LGS} = 2.3$ (for classification), while the original values for ILP from [2] are $t = 0.25$, and $s_{ILP} = 10$ (for optical-flow-based action recognition).

A.2. Defense-aware attack setup and parameters

Next, we evaluate the best combination of optimizers, learning rates (LR), and box constraints to optimize

¹Implementation from github.com/sniklaus/pytorch-spynet.




Figure A1. LGS and ILP hyperparameter study based on FlowNetC. Good parameters should balance the robustness against the vanilla attack $EPE(f, f_D^{Van})$ (dark color = good robustness) and small flow perturbations through the defense on unattacked frames $EPE(f, f_D)$ (dark color = small perturbation). We select $K = 16$, $O = 8$, $t = 0.15$, $b_{LGS} = 15$ and $s_{ILP} = 15$ as best parameters

defense-aware patches. As optimizers, we consider I-FGSM [21] and SGD, as learning rates 1, 0.1, 0.01 for I-FGSM and 10, 100 for SGD, and as box constraints either clipping or a change of variables (CoV). Due to the algorithmic differences between I-FGSM and SGD, the considered learning rates have distinct magnitudes to achieve comparable results. For each defense (none, LGS and ILP) we evaluate the pipeline robustness of the defended method under four separately trained defense-aware patches (using four fixed random seeds). We report the averaged robustness values for our defense-aware patches in Tab. A1 (no defense), Tab. A2 (LGS defense) and Tab. A3 (ILP defense).

Each defense-aware patch is trained for 2500 steps. The patches are randomly placed on the image, randomly rotated in a range of $[-10, 10]$ degrees, and randomly scaled in a range of $[0.95, 1.05]$. Batch size is chosen as 1 as the

Table A1. Robustness $EPE(f, f^{\text{Van}})$ [39] for undefended networks under vanilla patch attacks with different optimization parameter combinations. Non-evaluated settings are marked by “n.e.”.

Optim.	LR	Box	FNC	FNCR	PWC	SpyNet	RAFT	GMA	FF
SGD	10.00	CoV	61.86	0.73	1.34	1.10	0.27	0.29	0.42
SGD	10.00	Clip	52.41	0.97	1.17	1.01	0.28	0.31	0.45
SGD	100.00	CoV	76.28	0.62	1.28	1.26	0.29	0.34	n.e.
SGD	100.00	Clip	63.74	0.44	1.17	1.26	0.27	0.28	0.49
IFGSM	0.01	CoV	58.56	1.28	1.84	1.30	0.29	0.61	n.e.
IFGSM	0.01	Clip	32.19	1.58	1.80	1.19	0.29	0.55	0.54
IFGSM	0.10	CoV	57.55	1.47	1.84	1.33	0.34	0.46	n.e.
IFGSM	0.10	Clip	55.92	0.50	1.03	1.15	0.24	0.30	0.49
IFGSM	1.00	CoV	60.62	1.23	1.60	1.33	0.34	0.41	n.e.
IFGSM	1.00	Clip	8.22	0.45	0.88	1.11	0.26	0.32	0.44

Table A2. Robustness $EPE(f_{\text{LGS}}, f_{\text{LGS}}^{\text{LGS}})$ for LGS-defended networks under LGS-aware patch attacks with different optimization parameter combinations. Non-evaluated settings are marked by “n.e.”, while diverging optimization runs are marked as “div”.

Optim.	LR	Box	FNC	FNCR	PWC	SpyNet	RAFT	GMA	FF
SGD	10.00	CoV	3.98	3.07	3.28	3.62	1.45	1.57	n.e.
SGD	10.00	Clip	3.02	2.98	3.03	3.50	1.31	1.49	1.59
SGD	100.00	CoV	3.44	3.17	3.29	3.74	1.47	div.	n.e.
SGD	100.00	Clip	3.27	3.17	3.15	3.59	1.33	div.	1.56
IFGSM	0.01	CoV	22.64	2.51	3.74	3.69	1.05	1.24	n.e.
IFGSM	0.01	Clip	19.03	2.71	3.90	3.67	1.17	1.34	1.31
IFGSM	0.10	CoV	20.70	2.89	3.68	3.98	1.33	1.49	n.e.
IFGSM	0.10	Clip	8.42	2.62	3.04	3.60	1.22	1.33	1.28
IFGSM	1.00	CoV	4.61	3.10	3.28	3.62	1.42	1.53	n.e.
IFGSM	1.00	Clip	3.48	3.28	3.37	3.86	1.45	1.62	1.71

Table A3. Robustness $EPE(f_{\text{ILP}}, f_{\text{ILP}}^{\text{ILP}})$ for ILP-defended networks under ILP-aware patch attacks with different optimization parameter combinations. Non-evaluated settings are marked by “n.e.”, while diverging optimization runs are marked as “div”.

Optim.	LR	Box	FNC	FNCR	PWC	SpyNet	RAFT	GMA	FF
SGD	10.00	CoV	11.55	1.53	2.21	1.73	1.40	1.46	n.e.
SGD	10.00	Clip	4.09	2.99	2.99	2.17	1.45	1.48	1.75
SGD	100.00	CoV	3.17	2.90	3.08	2.52	1.43	div.	n.e.
SGD	100.00	Clip	3.56	3.27	3.37	2.83	1.42	1.55	1.69
IFGSM	0.01	CoV	57.46	2.95	3.84	2.77	1.12	1.25	n.e.
IFGSM	0.01	Clip	42.87	2.91	3.87	2.72	1.08	1.24	1.15
IFGSM	0.10	CoV	54.74	3.30	3.87	3.15	1.36	1.46	n.e.
IFGSM	0.10	Clip	18.70	2.93	3.11	2.77	1.23	1.32	1.39
IFGSM	1.00	CoV	3.84	3.22	3.34	2.98	1.37	1.43	n.e.
IFGSM	1.00	Clip	3.58	3.28	3.42	2.78	1.48	1.54	1.82

effect of batch size on the patch training is negligible. Due to its size and the resulting computational cost to evaluate

Table A4. Optimal parameter setups for defense-aware patch attacks on all optical flow methods with defenses. LR is the learning rate, and Box indicates whether a change of variables or clipping is used during optimization. The settings are a summary of the best results from Tab. A1, Tab. A2 and Tab. A3.

Attacked model	Defense	Attack	Optimizer	LR	Constraint
FlowNetC	None	Vanilla	SGD	100.00	CoV
	LGS	+LGS	IFGSM	0.01	CoV
	ILP	+ILP	IFGSM	0.01	CoV
FNCR	None	Vanilla	IFGSM	0.01	Clip
	LGS	+LGS	IFGSM	1.00	Clip
	ILP	+ILP	IFGSM	0.10	CoV
SpyNet	None	Vanilla	IFGSM	0.10	CoV
	LGS	+LGS	IFGSM	0.10	CoV
	ILP	+ILP	IFGSM	0.10	CoV
PWCNet	None	Vanilla	IFGSM	0.01	CoV
	LGS	+LGS	IFGSM	0.01	Clip
	ILP	+ILP	IFGSM	0.01	Clip
RAFT	None	Vanilla	IFGSM	1.00	CoV
	LGS	+LGS	SGD	100.00	CoV
	ILP	+ILP	IFGSM	1.00	Clip
GMA	None	Vanilla	IFGSM	0.01	CoV
	LGS	+LGS	IFGSM	1.00	Clip
	ILP	+ILP	IFGSM	1.00	Clip
FlowFormer	None	Vanilla	IFGSM	0.01	Clip
	LGS	+LGS	IFGSM	1.00	Clip
	ILP	+ILP	IFGSM	1.00	Clip

FlowFormer [15], we only train its patches for 1000 iterations and omit the change of variables to reduce the number of test runs, as it performed similarly to RAFT and GMA. We found patches to be sufficiently converged after the reduced number of iterations. Please note that across all methods, the choice of box constraint did not significantly influence the effectiveness of the adversarial patches. The patch optimization for GMA diverged with SGD and learning rate 100 for LGS- and ILP-aware patches.

Based on this extensive parameter evaluation, we select the best optimization parameters for all combinations of optical flow network and defense-aware attack in Tab. A4, which are boldfaced in the detailed evaluations in Tab. A1, Tab. A2 and Tab. A3. These parameters were used to produce the defense-aware patches for the experimental evaluation in the Main paper.

Additionally, we show the best (out of four) defense-aware patches for no defense, LGS-defense and ILP-defense in Fig. A2, Fig. A3 and Fig. A4, respectively. The best patch is selected based on the worst robustness score of the defended method after training.

A.3. Additional flow visualizations for vanilla attack

Here, we complement the limited selection of methods whose optical flow was visualized for unattacked and (vanilla) attacked frames in Main Fig. 4. Unattacked and vanilla-attacked flow visualizations on *all* tested optical

Optim.	LR	Box	FlowNetC	FlowNetCRobust	PWCNet	SpyNet	RAFT	GMA	FlowFormer
SGD	10.00	CoV							
SGD	10.00	Clip							
SGD	100.00	CoV							n.e.
SGD	100.00	Clip							
IFGSM	0.01	CoV							n.e.
IFGSM	0.01	Clip							
IFGSM	0.10	CoV							n.e.
IFGSM	0.10	Clip							
IFGSM	1.00	CoV							n.e.
IFGSM	1.00	Clip							

Figure A2. Best-performing vanilla patches for different networks and optimization parameter combinations. Non-evaluated settings are marked by “n.e.”. See Tab. A1 for the corresponding robustness values, averaged over four patches.

flow methods for the previous KITTI scene are in Fig. A5 and for an additional KITTI sample in Fig. A6. For a lean representation, only a single frame of the attacked image pair is shown on the right.

In both figures, RAFT [44], GMA [18] and FlowFormer [15] are able to recognize the patch as a static object

in the scene and therefore predict its output flow as zero. The less accurate methods SpyNet [34], PWCNet [42] and FlowNetCRobust [40] also recognize the zero flow, but their flow predictions are overall less precise and the patch bleeds into the surrounding area. The outlier is FlowNetC [11], where the entire flow prediction is deteriorated by the patch.

Optim.	LR	Box	FlowNetC	FlowNetCRobust	PWCNet	SpyNet	RAFT	GMA	FlowFormer
SGD	10.00	CoV							n.e.
SGD	10.00	Clip							
SGD	100.00	CoV						div.	n.e.
SGD	100.00	Clip						div.	
IFGSM	0.01	CoV							n.e.
IFGSM	0.01	Clip							
IFGSM	0.10	CoV							n.e.
IFGSM	0.10	Clip							
IFGSM	1.00	CoV							n.e.
IFGSM	1.00	Clip							

Figure A3. Best-performing LGS-aware patches for different networks and optimization parameter combinations. Non-evaluated settings are marked by “n.e.,” while diverging optimization runs are marked as “div.” See Tab. A2 for the corresponding robustness values, averaged over four patches.

In both visualizations, almost all optical flow methods are hardly affected by the patch, as they correctly recognize it as an object and accurately predict its zero motion.

A.4. Manual patch attack: Defended quality

In the manual patch analysis in Sec. 6.4, the Main paper visually argued that our high-frequent, manual patch attacks qualitatively improve the optical flow predictions of LGS- and ILP-defended methods, as a result of the sig-

Optim.	LR	Box	FlowNetC	FlowNetCRobust	PWCNet	SpyNet	RAFT	GMA	FlowFormer
SGD	10.00	CoV							n.e.
SGD	10.00	Clip							
SGD	100.00	CoV						div.	n.e.
SGD	100.00	Clip							
IFGSM	0.01	CoV							n.e.
IFGSM	0.01	Clip							
IFGSM	0.10	CoV							n.e.
IFGSM	0.10	Clip							
IFGSM	1.00	CoV							n.e.
IFGSM	1.00	Clip							

Figure A4. Best-performing vanilla patches for different networks and optimization parameter combinations. Non-evaluated settings are marked by “n.e.,” while diverging optimization runs are marked as “div”. See Tab. A3 for the corresponding robustness values, averaged over four patches.

nificant quality degradation of defenses on unattacked images. Here, we provide the corresponding quality scores over the whole set of KITTI frames. To this end, we quantify the quality $Q_D^A = \text{EPE}(f^*, f_D^A)$, *i.e.* the distance between ground truth flow f^* and optical flow predictions f_D^A of methods that are defended with D and attacked with A.

Tab. A5 provides the quality scores for unattacked but defended networks (block 1, corresponds to values from Main Tab. 1), for full pipelines where the defended method is attacked with the corresponding defense-awareness (block 2) and for our manual attack on defended networks (block 3). Compared to the original baseline

Table A5. Quality $Q_D^A = \text{EPE}(f^*, f_D^A)$, *i.e.* the distance between ground truth flow and optical flow predictions that are defended with D and attacked with A. The upper block shows the quality scores from Main Tab. 1 (for comparison), and the lower blocks contain the quality scores for full pipelines and manual patch attacks on networks with varying defenses. Per method, we mark the best defended quality for unattacked networks and full pipelines **bold** (includes the first two blocks, up to double line), and underline the best quality if the manual patch is also included – the undefended baselines that are marked in gray are excluded from both rankings.

u Attack type	Defense		FNC	FNCR	PWC	SpyNet	RAFT	GMA	FF
No Attack	None	Q	15.42	11.10	13.26	24.03	0.63	0.61	0.62
	LGS	Q_{LGS}	16.70	13.13	14.61	25.15	1.42	1.55	1.42
	ILP	Q_{ILP}	16.46	12.77	14.52	24.74	1.36	1.39	1.30
Vanilla	None	Q^{Van}	84.48	12.64	15.27	25.11	0.80	0.91	0.78
+LGS (LGS-aware)	LGS	$Q_{\text{LGS}}^{\text{LGS}}$	34.41	11.68	15.43	25.28	0.94	0.90	0.83
+ILP (ILP-aware)	ILP	$Q_{\text{ILP}}^{\text{ILP}}$	65.02	12.31	15.65	25.29	0.68	0.70	0.68
Manual	None	Q^{Man}	16.21	11.38	13.87	24.79	0.71	0.70	0.70
	LGS	$Q_{\text{LGS}}^{\text{Man}}$	16.66	11.69	14.24	24.87	0.92	0.91	0.88
	ILP	$Q_{\text{ILP}}^{\text{Man}}$	<u>16.16</u>	<u>11.52</u>	<u>13.97</u>	<u>24.69</u>	0.70	<u>0.68</u>	0.72

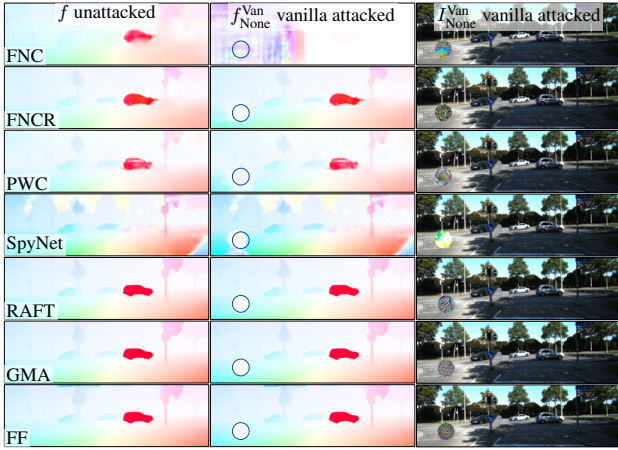


Figure A5. Unattacked optical flow estimation (left) and corresponding vanilla-attacked optical flow (middle) for all tested methods on a KITTI sample (right). Complements Main Fig. 4, see Fig. A6 for more samples.

$Q = \text{EPE}(f^*, f)$ (block 1, marked in gray), all defenses decrease the quality. Fig. A7 and Fig. A8 visualize the flow output for unattacked models on KITTI samples when no defense, LGS or ILP is applied.

Then, we begin by comparing the quality for defended methods in the first two blocks, *i.e.* we exclude the manual patch attack in block 3, and mark the best quality **bold** in the table. Note that we exclude the gray rows, as they contain the quality for *undefended* methods. For optical flow methods that have a good undefended quality Q , *i.e.* FlowNetCRobust, RAFT, GMA and FlowFormer, we find that a defense-aware attack on a defended model actually yields a better quality than the defended but unattacked

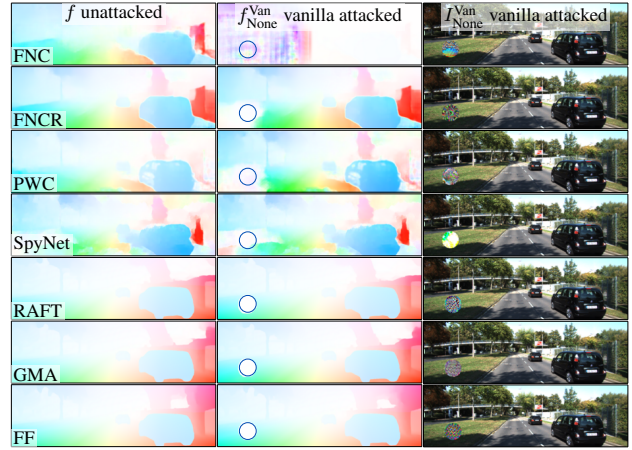


Figure A6. Unattacked optical flow estimation (left) and corresponding vanilla-attacked optical flow (middle) for all tested methods on a KITTI sample (right). See Fig. A5 for more samples.

model: $Q_D^D > Q_D$. For these methods, a noisy patch was revealed to be the most effective. Hence, it is easier for an adaptive attack to exploit the changes introduced by the defense than to influence the flow estimation.

Now we also include the manual patch attack in the defense evaluation, again underlining the highest-quality flow per method over all three blocks in Tab. A5. Again we exclude the gray rows that contain the quality for *undefended* methods in order to compare the influence of the defenses. Now, for almost all methods the best defended quality is achieved for manual patch attacks. When we compare the underlined numbers to the baseline quality Q , we find that our manual patch attack almost restores the undefended and unattacked quality for our defended methods outside the

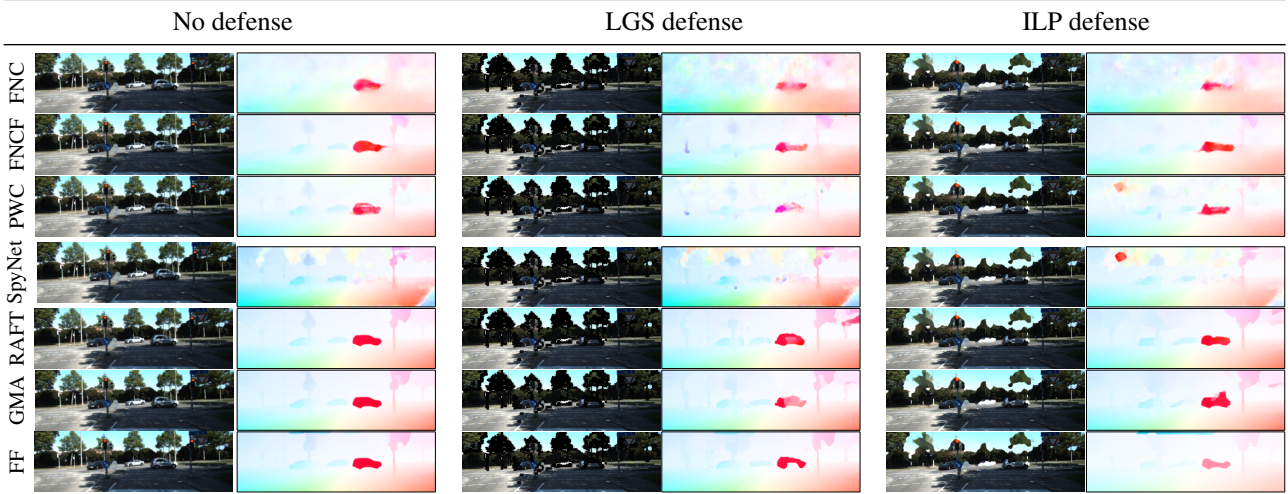


Figure A7. Optical flow prediction on an unattacked frame of the KITTI dataset for optical flow methods with different defenses. Defenses from left to right: None, LGS and ILP. See Fig. A8 for more samples.

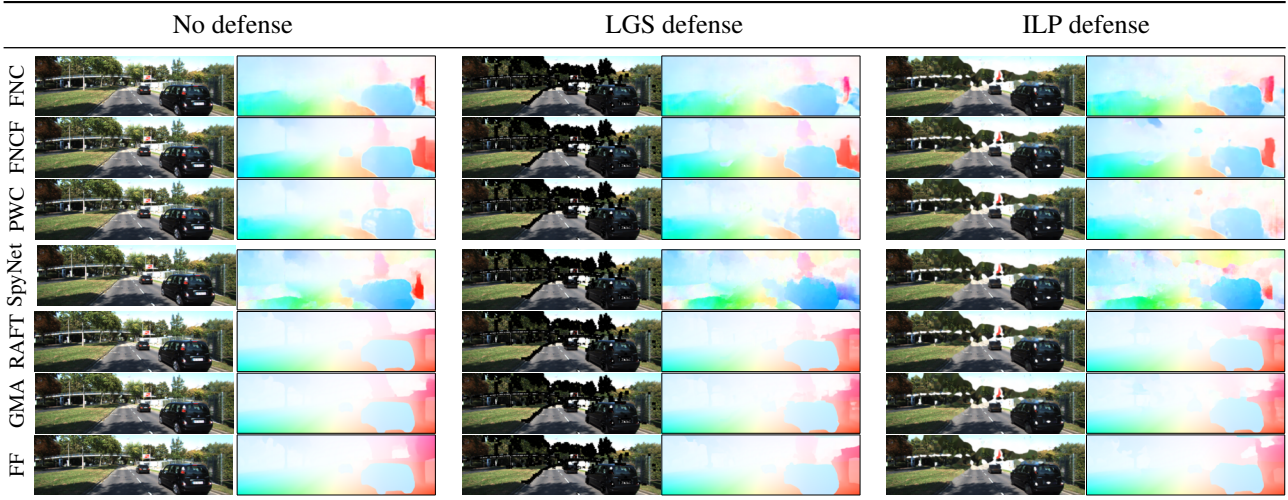


Figure A8. Optical flow prediction on an unattacked frame of the KITTI dataset for optical flow methods with different defenses. Defenses from left to right: None, LGS and ILP. See Fig. A7 for more samples.

patch area. While this underlines the finding from the Main paper that the low quality of defended but unattacked methods is the main reason for the low quality (and robustness) of defended methods, it also yields another point: If the defenses did not deteriorate the unattacked quality, they could be effective in terms of quality and robustness because they restore high-quality optical flow fields in the presence of adversarial-like patches.

A.5. Defense evaluation on additional datasets

We evaluate the defenses and their effectiveness on more datasets besides KITTI [28], and consider Sintel [6], Driving [25], HD1K [19] and Spring [27]. Because evaluating the defended quality requires ground truth optical flow

data, we use validation splits of the respective test sets for all datasets. Dataset-specific patches are then trained on the remaining training data. For Sintel, we use the validation set from [51] which splits Sintel-test such that the flow magnitudes of the validation set match the flow-magnitude distribution of the full training set [41]. For HD1K and Spring, we are unaware of flow-magnitude matching validation splits in the literature, and create validation splits with matching flow-magnitude distributions as detailed in Tab. A6. For Driving, we use the scenes with focal length 15mm, forwards, fast speed and left camera as validation split. Note that during our evaluations, we half the image resolution for HD1K and Spring, to keep the image sizes and hence results for patches with size 100 compa-

Table A6. Validation split details for the evaluation datasets. “Frames” denotes frame pairs (for the optical flow calculation) rather than single frames, if “Half” is checked the frame size is halved. If all scenes except the validation scenes make up the set for training patches, the training scenes are marked with “EV: except validation”. “OFM-id” denotes “optical flow magnitude in-distribution”, meaning the validation set is in-distribution w.r.t. to the optical flow magnitude distribution of the original training set. “I3” means that only every third frame pair of the validation scenes is added to the validation split.

Dataset	Val. Frames	Val. scenes	Half	Train. scenes	Notes
KITTI [28]	200	KITTI-train [28]	–	Raw [12]	Split from [35]
Sintel [6]	89	ambush2, bamboo2, cave2, market2, shaman2, temple2	–	EV	Split from [51] OFM-id, I3
Driving [25]	299	15mm focal length, scene forwards, fast, into future, left	–	EV	
HD1K [19]	94	000009, 000013, 000018, 000019, 000032	✓	EV	OFM-id
Spring [27]	658	0002, 0010, 0018, 0026, 0032, 0045	✓	EV	OFM-id

Table A7. Quality $Q_D = EPE(f^*, f_D)$ for optical flow pipelines with defense D on the HD1K [19] validation split²; Best quality is **bold**. All defenses lead to a worse quality on unattacked frames.

Defense	FNC	FNCR	SpyNet	PWC	RAFT	GMA	FF
None Q	2.36	1.17	3.00	2.15	0.46	0.44	0.32
LGS Q_{LGS}	2.49	1.22	3.09	2.19	0.50	0.47	0.35
ILP Q_{ILP}	2.39	1.20	3.10	2.21	0.50	0.46	0.35

table across all datasets.

For the datasets HD1K, Spring, Sintel (clean and final) and Driving (clean and final), we show the numerical results of the defended quality analysis in Tab. A7, Tab. A8, Tab. A9 and Tab. A10, and the respective robustness analyses in Tab. A11, Tab. A12, Tab. A13 and Tab. A14. For a better overview, Fig. A9 shows the quality vs. robustness plots for all tested optical flow methods on all tested datasets, which can be compared to the results on KITTI in Main Fig. 5.

Focusing on the quality-robustness plots in Fig. A9, we observe that defenses worsen quality and robustness for all optical flow methods (except those of FlowNetC) on HD1K and Spring, cf. Fig. A9a and Fig. A9d. On Sintel and Driving, the results are more differentiated: For high-quality methods like RAFT, GMA and FlowFormer (red markers), defending them with ILP improves the robustness for the

²See Tab. A6 for details on the used validation splits

Table A8. Quality $Q_D = EPE(f^*, f_D)$ for optical flow pipelines with defense D on the Spring [27] validation split²; Best quality is **bold**. All defenses lead to a worse quality on unattacked frames.

Defense	FNC	FNCR	SpyNet	PWC	RAFT	GMA	FF
None Q	0.81	0.48	0.96	1.87	0.29	0.29	0.27
LGS Q_{LGS}	1.17	1.09	1.39	2.21	0.66	0.67	0.44
ILP Q_{ILP}	1.10	0.74	1.21	2.07	0.40	0.48	0.39

Table A9. Quality $Q_D = EPE(f^*, f_D)$ for optical flow pipelines with defense D on the Sintel [6] clean and final validation splits² [51]; Best quality is **bold**. All defenses lead to a worse quality on unattacked frames.

Defense	FNC	FNCR	SpyNet	PWC	RAFT	GMA	FF
clean							
None Q	4.80	2.31	5.66	3.72	0.84	0.77	0.45
LGS Q_{LGS}	4.83	2.34	5.81	3.75	0.86	0.79	0.48
ILP Q_{ILP}	4.90	2.37	5.70	3.77	0.85	0.80	0.49
final							
None Q	5.79	4.07	7.64	5.42	1.49	1.45	0.74
LGS Q_{LGS}	5.82	4.16	7.66	5.51	1.56	1.51	0.79
ILP Q_{ILP}	5.81	4.10	7.63	5.46	1.52	1.47	0.76

Table A10. Quality $Q_D = EPE(f^*, f_D)$ for optical flow pipelines with defense D on the Driving [25] clean and final validation splits²; Best quality is **bold**. All defenses lead to a worse quality on unattacked frames.

Defense	FNC	FNCR	SpyNet	PWC	RAFT	GMA	FF
clean							
None Q	95.73	88.85	111.36	92.49	37.32	54.17	51.94
LGS Q_{LGS}	96.20	89.64	113.15	92.93	36.63	60.21	53.11
ILP Q_{ILP}	95.99	89.43	112.32	93.06	36.75	60.06	53.21
final							
None Q	90.21	84.34	110.52	92.19	40.87	62.35	47.05
LGS Q_{LGS}	90.78	85.07	111.84	92.70	41.96	62.68	47.52
ILP Q_{ILP}	90.47	84.59	110.91	92.52	41.60	62.61	47.56

final versions of the datasets in Fig. A9e and Fig. A9f – on the clean dataset versions in Fig. A9b and Fig. A9c, however, both defenses deteriorate either quality, or robustness, or both. Defending the lower-quality methods SpyNet and PWCNet (blue markers) also deteriorates at least quality or robustness on both datasets, with the exception of PWCNet, where defending leads to minor robustness improvements on Sintel. For FlowNetC and FlowNetCRobust

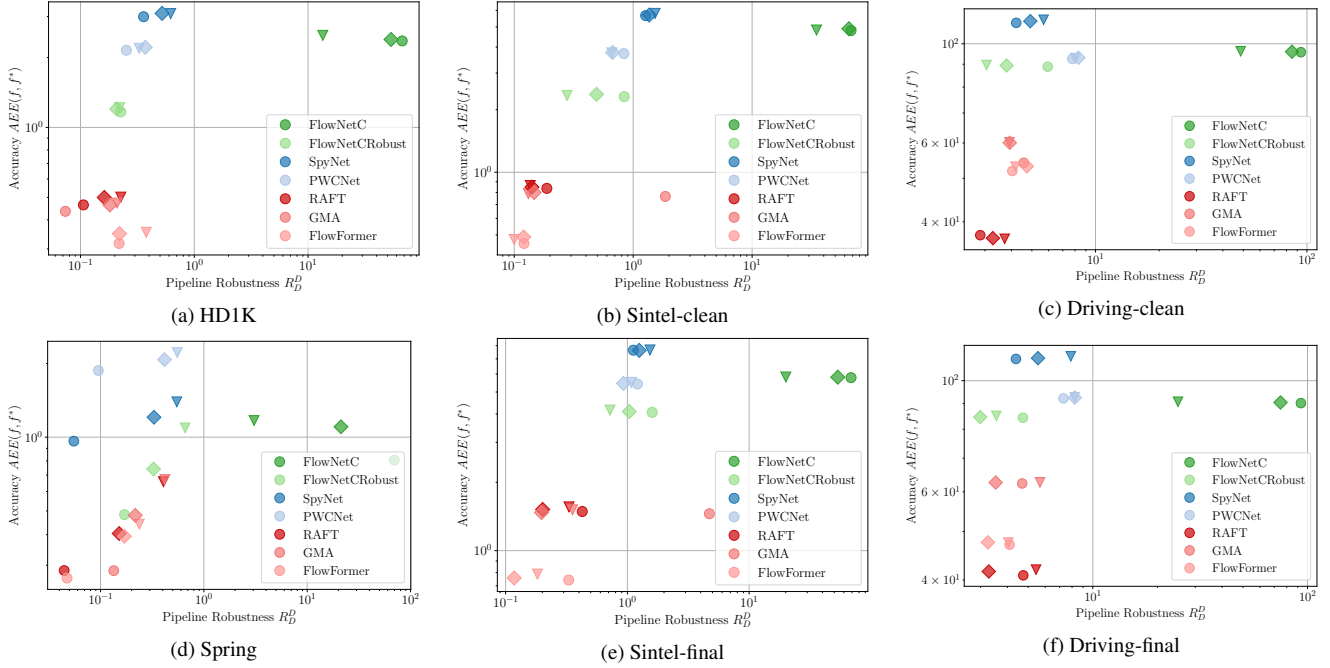


Figure A9. Quality vs. robustness of flow networks on different datasets in a double logarithmic plot. An ideal method would be in the origin. Undefended networks are circles \circ , networks defended with LGS are triangles ∇ and networks defended with ILP are diamonds \diamond .

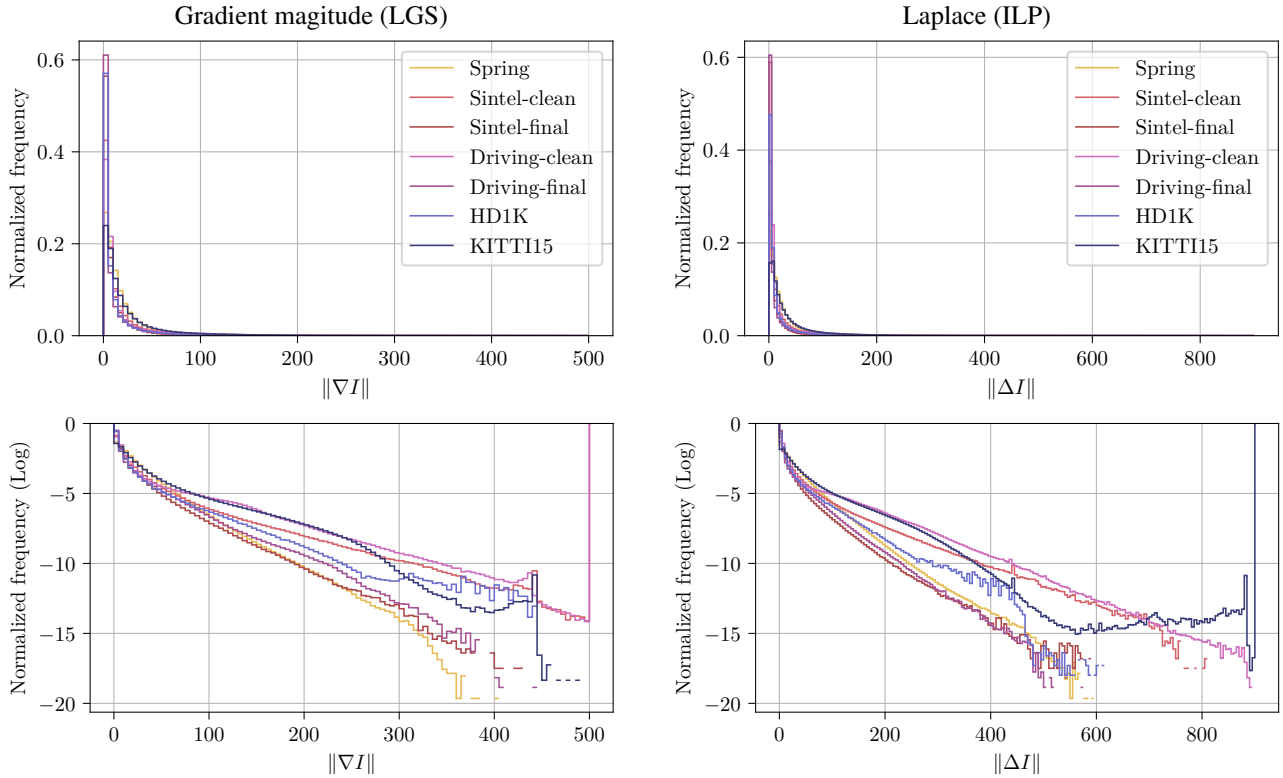


Figure A10. Image statistics for optical flow datasets. The plots show the histograms over the magnitude of first and second image derivatives for different optical flow datasets, where the LGS defense considers first (left) and ILP considers second (right) derivatives. The histograms are normalized by the number of pixels in the respective dataset. The top row shows the pure histograms, while the bottom row shows the log-transformed frequency for better visualization of statistics for large gradient magnitudes, which are filtered by the defenses.

Table A11. Robustness scores for all combinations of defended methods and defense-aware attacks on optical flow methods on the HD1K [19] validation split². Per attack, the robustness values of the best defense are **bold**. Per defense, the robustness values for the attack it is most vulnerable to are underlined. Full pipelines are highlighted in gray, and provide the corresponding robustness values to the quality scores from Tab. A7.

Attack type	Defense		FNC	FNCR	SpyNet	PWC	RAFT	GMA	FF
Vanilla	None	R_{Van}^{Van}	67.24	<u>0.23</u>	0.36	<u>0.25</u>	0.11	0.07	0.22
	LGS	R_{LGS}^{Van}	0.45	0.21	0.59	0.25	<u>0.23</u>	0.17	0.36
	ILP	R_{ILP}^{Van}	0.60	0.17	0.51	0.24	0.17	0.13	<u>0.22</u>
+LGS (LGS-aware)	None	R_{LGS}^{LGS}	51.97	0.07	0.35	0.22	0.09	0.06	0.17
	LGS	R_{LGS}^{LGS}	<u>13.47</u>	<u>0.22</u>	<u>0.62</u>	<u>0.33</u>	0.23	0.21	0.38
	ILP	R_{ILP}^{LGS}	10.02	0.17	0.51	0.35	<u>0.18</u>	<u>0.18</u>	0.22
+ILP (ILP-aware)	None	R_{ILP}^{ILP}	60.52	0.14	0.36	0.24	0.06	0.06	0.18
	LGS	R_{LGS}^{ILP}	2.89	0.21	0.62	0.27	0.23	<u>0.21</u>	<u>0.38</u>
	ILP	R_{ILP}^{ILP}	<u>53.63</u>	<u>0.21</u>	<u>0.52</u>	<u>0.37</u>	0.16	0.18	0.22

Table A12. Robustness scores for all combinations of defended methods and defense-aware attacks on optical flow methods on the Spring [27] validation split². Per attack, the robustness values of the best defense are **bold**. Per defense, the robustness values for the attack it is most vulnerable to are underlined. Full pipelines are highlighted in gray, and provide the corresponding robustness values to the quality scores from Tab. A8.

Attack type	Defense		FNC	FNCR	SpyNet	PWC	RAFT	GMA	FF
Vanilla	None	R_{Van}^{Van}	69.64	<u>0.17</u>	0.06	0.10	0.04	0.13	0.05
	LGS	R_{LGS}^{Van}	0.62	0.67	0.55	0.54	0.40	<u>0.42</u>	<u>0.24</u>
	ILP	R_{ILP}^{Van}	0.58	0.31	0.33	0.36	0.15	<u>0.22</u>	<u>0.17</u>
+LGS (LGS-aware)	None	R_{LGS}^{LGS}	33.27	0.02	0.06	0.16	0.02	0.02	0.03
	LGS	R_{LGS}^{LGS}	3.06	0.66	<u>0.55</u>	<u>0.55</u>	<u>0.41</u>	0.42	0.24
	ILP	R_{ILP}^{LGS}	11.87	0.32	<u>0.33</u>	<u>0.42</u>	<u>0.15</u>	0.22	0.17
+ILP (ILP-aware)	None	R_{ILP}^{ILP}	30.01	0.10	0.06	<u>0.19</u>	0.01	0.02	0.03
	LGS	R_{LGS}^{ILP}	0.74	<u>0.67</u>	0.55	0.55	0.41	0.42	0.24
	ILP	R_{ILP}^{ILP}	<u>21.22</u>	<u>0.33</u>	0.33	0.42	0.15	0.22	0.17

(green markers), defenses do indeed improve the robustness on Sintel and Driving, but here it is LGS that leads to the best robustness scores. Overall, this clearly supports that defenses should not be used in a “plug’n’play” manner without extensive application-specific testing, as they either do not improve the optical flow methods at all, or – when they do improve the robustness – their effect is small and does not apply to more than a few selected optical flow methods. Hence, current detect-and-remove defenses cannot be recommended for general use.

To better understand the effectiveness differences of defenses on the tested datasets, we analyze the results in relation to the datasets in more detail. When we consider the datasets KITTI, HD1K and Spring and their quality-robustness plots in Main Fig. 5, Fig. A9a and Fig. A9d, we find that applying defenses to optical flow methods worsens quality *and* robustness, which leads to a slanted line of markers per optical flow network. This indicates that for these datasets, the defenses affect the unattacked de-

fended flow f_D as described in Sec. 6.4, Main paper, because worsening this flow enters into both, the quality calculation with $EPE(f^*, f_D)$ and the robustness calculation with $EPE_P(f_D, f_D^A)$. For the datasets Sintel and Driving in Fig. A9e, Fig. A9b, Fig. A9f and Fig. A9c, applying defenses almost exclusively changes the robustness, leading to a horizontal line of markers per optical flow network. This indicates that the defenses work “as intended”, affecting only the attacked defended flow f_D^A and hence the robustness, but are on average not very effective under attack with defense-aware patches. In summary, defenses have the worst side effects on the image quality for natural or naturalistic data: KITTI and HD1K contain camera-captured real-world images and Spring is a recently rendered dataset that focuses on high-detail images. Even though defenses work partially on the synthetic datasets Sintel and Driving, which were rendered and created before 2016, they still fail to demonstrate consistent advantages over undefended networks on these datasets. These differences in the datasets

Table A13. Robustness scores for all combinations of defended methods and defense-aware attacks on optical flow methods on the Sintel [6] final (f) and clean (c) validation splits² [51]. Per attack, the robustness values of the best defense are **bold**. Per defense, the robustness values for the attack it is most vulnerable to are underlined. Full pipelines are highlighted in gray, and provide the corresponding robustness values to the quality scores from Tab. A9.

Attack type	Defense		FNC		FNCR		SpyNet		PWC		RAFT		GMA		FF	
			f	c	f	c	f	c	f	c	f	c	f	c	f	c
Vanilla	None	R^{Van}	68.71	67.84	1.60	0.84	1.12	1.27	1.21	0.84	0.43	0.19	4.72	1.86	0.33	0.12
	LGS	$R_{\text{LGS}}^{\text{Van}}$	1.06	0.90	0.73	0.25	1.45	1.41	0.86	0.52	0.40	0.17	0.34	0.13	0.19	0.09
	ILP	$R_{\text{ILP}}^{\text{Van}}$	1.49	1.76	0.63	0.29	1.16	1.34	0.74	0.51	0.19	0.14	0.24	0.13	0.13	<u>0.12</u>
+LGS (LGS-aware)	None	R^{LGS}	57.21	57.11	0.44	0.20	1.15	1.23	0.90	0.61	0.23	0.10	0.20	0.09	0.16	0.07
	LGS	$R_{\text{LGS}}^{\text{LGS}}$	20.04	34.78	0.72	<u>0.28</u>	1.53	1.53	1.08	0.66	0.33	0.14	0.35	0.13	0.18	0.10
	ILP	$R_{\text{ILP}}^{\text{LGS}}$	31.42	40.56	0.63	0.35	1.23	1.33	0.88	0.64	0.20	0.13	0.19	0.15	0.12	0.12
+ILP (ILP-aware)	None	R^{ILP}	68.58	<u>68.13</u>	1.13	0.55	1.17	1.26	0.96	0.65	0.20	0.09	0.19	0.09	0.19	0.08
	LGS	$R_{\text{LGS}}^{\text{ILP}}$	2.54	24.19	0.74	0.27	1.53	1.53	0.98	0.64	0.35	0.16	0.35	0.13	0.19	<u>0.10</u>
	ILP	$R_{\text{ILP}}^{\text{ILP}}$	<u>53.52</u>	<u>65.18</u>	1.04	<u>0.49</u>	1.25	1.36	0.93	0.67	0.20	0.14	0.20	<u>0.15</u>	0.12	0.12

Table A14. Robustness scores for all combinations of defended methods and defense-aware attacks on optical flow methods on the Driving [25] final (f) and clean (c) validation splits². Per attack, the robustness values of the best defense are **bold**. Per defense, the robustness values for the attack it is most vulnerable to are underlined. Full pipelines are highlighted in gray, and provide the corresponding robustness values to the quality scores from Tab. A10.

Attack type	Defense		FNC		FNCR		SpyNet		PWC		RAFT		GMA		FF	
			f	c	f	c	f	c	f	c	f	c	f	c	f	c
Vanilla	None	R^{Van}	93.15	3.64	4.72	5.94	4.39	4.22	7.30	7.77	4.74	2.85	4.69	4.58	4.10	4.04
	LGS	$R_{\text{LGS}}^{\text{Van}}$	4.43	3.25	3.39	2.12	7.72	5.33	6.83	6.69	5.49	3.41	5.18	3.40	3.91	3.44
	ILP	$R_{\text{ILP}}^{\text{Van}}$	4.88	<u>6.05</u>	2.54	2.23	5.54	4.90	5.73	7.33	3.31	2.83	3.43	3.19	3.26	3.38
+LGS (LGS-aware)	None	R^{LGS}	74.11	4.11	1.30	1.64	4.25	4.07	6.63	7.00	4.19	2.90	3.14	3.01	3.53	3.71
	LGS	$R_{\text{LGS}}^{\text{LGS}}$	<u>24.93</u>	8.61	<u>3.55</u>	<u>3.05</u>	7.90	5.68	8.24	<u>8.02</u>	5.44	<u>3.72</u>	<u>5.67</u>	3.91	4.05	4.18
	ILP	$R_{\text{ILP}}^{\text{LGS}}$	24.21	5.93	2.48	<u>4.07</u>	5.49	5.18	7.96	7.90	3.98	<u>3.77</u>	3.51	3.87	3.16	4.54
+ILP (ILP-aware)	None	R^{ILP}	84.98	5.34	2.70	3.49	4.34	4.17	6.82	7.44	3.05	2.56	3.15	2.96	3.74	3.93
	LGS	$R_{\text{LGS}}^{\text{ILP}}$	7.92	8.20	3.32	2.38	7.86	<u>5.68</u>	7.25	7.58	5.46	3.69	5.52	<u>4.04</u>	4.16	<u>4.42</u>
	ILP	$R_{\text{ILP}}^{\text{ILP}}$	<u>74.86</u>	5.00	<u>2.99</u>	3.80	<u>5.55</u>	<u>4.91</u>	8.24	<u>8.30</u>	3.28	3.27	<u>3.53</u>	<u>3.93</u>	3.25	4.73

are also visible in terms of the dataset image statistics that are considered by the LGS and ILP defenses. In Fig. A10 we show the histograms over first- and second-order image derivatives for all datasets. There, the synthetic Sintel and Driving datasets have a very “even” gradient magnitude decay for large gradients on the log scale (both for clean and final rendering passes), while the realistic KITTI, HD1K and Spring datasets do not show such a clear exponential gradient decay. All in all, defenses fail most severely on safety-critical real-world datasets, where reliable predictions are needed most.