

# 000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 ROBUSTSPRING: BENCHMARKING ROBUSTNESS TO IMAGE CORRUPTIONS FOR OPTICAL FLOW, SCENE FLOW AND STEREO

Anonymous authors

Paper under double-blind review

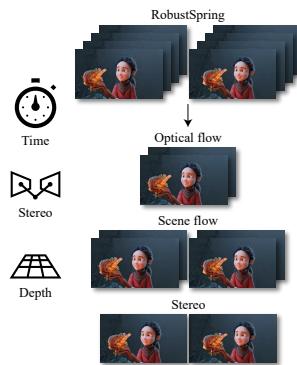
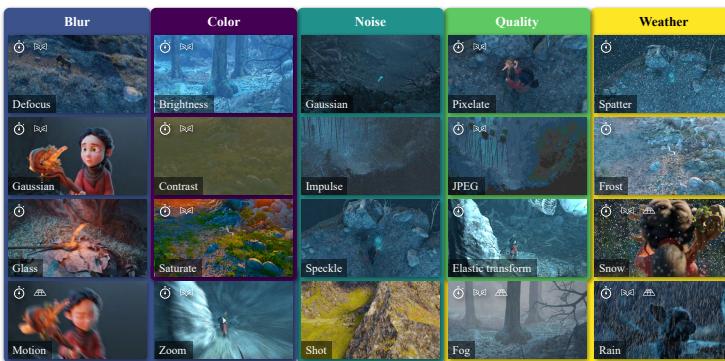


Figure 1: RobustSpring is a novel image corruption benchmark for optical flow, scene flow, and stereo. It evaluates 20 image corruptions including blurs, color changes, noises, quality degradations, and weather, applied to stereo video data from Mehl et al. (2023b). For comprehensive robustness evaluations on all three tasks, RobustSpring’s image corruptions are integrated in time, stereo, and depth, where applicable.

## ABSTRACT

Standard benchmarks for optical flow, scene flow, and stereo vision algorithms generally focus on model accuracy rather than robustness to image corruptions like noise or rain. Hence, the resilience of models to such real-world perturbations is largely unquantified. To address this, we present RobustSpring, a comprehensive dataset and benchmark for evaluating robustness to image corruptions for optical flow, scene flow, and stereo models. RobustSpring applies 20 different image corruptions, including noise, blur, color changes, quality degradations, and weather distortions, in a time-, stereo-, and depth-consistent manner to the high-resolution Spring dataset, creating a suite of 20,000 corrupted images that reflect challenging conditions. RobustSpring enables comparisons of model robustness via a new corruption robustness metric. Integration with the Spring benchmark enables two-axis evaluations of both accuracy and robustness. We benchmark a curated selection of initial models, observing that robustness varies widely by corruption type, and experimentally show that evaluations on RobustSpring indicate real-world robustness. RobustSpring is a new computer vision benchmark that treats robustness as a first-class citizen to foster models that combine accuracy with resilience.

## 1 INTRODUCTION

Optical flow, scene flow, and stereo vision algorithms estimate dense correspondences and enable real-world applications like robot navigation (McGuire et al., 2017; Zhang et al., 2025; Lamberti et al., 2024), video processing (Mehl et al., 2024), structure-from-motion (Maurer et al., 2018; Phan et al., 2020), medical image registration (Mocanu et al., 2021), or surgical assistance (Rosa et al., 2019; Philipp et al., 2022). While estimation quality continuously improves on accuracy-driven benchmarks (Mehl et al., 2023b; Menze & Geiger, 2015; Butler et al., 2012; Baker et al., 2011; Scharstein et al., 2014; Geiger et al., 2012; Richter et al., 2017; Schöps et al., 2017), their robustness

054 to real-world visual corruptions like noise or compression artifacts is rarely systematically assessed.  
 055 This lack of systematic assessment is problematic, as better accuracy does not necessarily trans-  
 056 late to improved robustness, and can even harm model robustness (Tsipras et al., 2019; Schmalfuss  
 057 et al., 2022b). Though image data in KITTI (Menze & Geiger, 2015), Sintel (Butler et al., 2012),  
 058 or Spring (Mehl et al., 2023b) comes with degradations like motion blurs, depth-of-field, or bright-  
 059 ness changes, they result from real-world data capture or efforts to increase data realism, but were  
 060 not included to systematically study model predictions under image corruptions. Broad corruption-  
 061 robustness studies as they exist for image classification (Hendrycks & Dietterich, 2019; Müller et al.,  
 062 2023), 3D object detection (Michaelis et al., 2019; Kong et al., 2023) or monocular depth estima-  
 063 tion (Kar et al., 2022) are rare for dense-correspondence tasks, where studies are limited to specific  
 064 degradations like weather (Schmalfuss et al., 2023) or low-light (Zheng et al., 2020). This not only  
 065 leaves uncertainty about the reliability of dense matching algorithms in real-world scenarios. It also  
 066 prevents systematic efforts to improve their robustness.  
 067

068 To enable systematic studies on the image corruption robustness of optical flow, scene flow, and  
 069 stereo, we propose the *RobustSpring* dataset. Based on Spring (Mehl et al., 2023b), it jointly bench-  
 070 marks robustness of all three tasks on corrupted stereo videos. While prior image corruptions affect  
 071 the monocular 2D or 3D space (Hendrycks & Dietterich, 2019; Kar et al., 2022; Michaelis et al.,  
 072 2019), RobustSpring’s image corruptions are integrated in *time*, *stereo*, and *depth* and thus tailored to  
 073 dense matching tasks. A principled corruption robustness metric and an accompanying benchmark  
 074 framework make RobustSpring the first systematic tool to evaluate and improve dense matching  
 075 robustness to image corruptions.  
 076

**Contributions.** Figure 1 gives an overview of RobustSpring. In summary, we make the following  
 077 contributions:

- (1) *Tailored image corruptions.* RobustSpring is the first image corruption dataset for optical  
 078 flow, scene flow, and stereo. It integrates 20 corruptions for blurs, noises, tints, artifacts,  
 079 and weather in time, stereo, and depth.
- (2) *Corruption robustness metric.* We propose a corruption robustness metric based on Lips-  
 080 chitz continuity, which subsamples the clean-corrupted prediction difference and disentan-  
 081 gles robustness and accuracy.
- (3) *Benchmark functionality.* RobustSpring’s standardized evaluation enables community-  
 082 driven robustness comparisons of dense matching models. Public robustness benchmarking  
 083 can be integrated with Spring’s website.
- (4) *Initial robustness evaluation.* We benchmark nine optical flow, two scene flow, and six  
 084 stereo models. All models are corruption sensitive, which reveals concealed robustness  
 085 deficits on dense matching models.

**Intended Use.** RobustSpring is not a fine-tuning dataset, but a benchmark of how dense matching  
 090 models generalize to *unseen* image corruptions. It seeks to foster robustness research and, simulta-  
 091 neously, helps assess real-world applicability of models. Hence, it is essential to tie RobustSpring  
 092 to an existing accuracy benchmark like Spring, as this minimizes the robustness evaluation hurdle  
 093 for researchers. While RobustSpring treats corruptions as perturbations to assess robustness, their  
 094 interpretation depends on the application domain. For instance, in autonomous driving, rain or snow  
 095 are typically considered disturbances to be ignored for stable navigation, whereas in video editing  
 096 or cinematic rendering, they may constitute meaningful scene content. To accommodate such differ-  
 097 ences, RobustSpring provides results per corruption type, allowing end-users to focus only on those  
 098 corruptions that align with their intended application.

## 100 2 RELATED WORK

101 While the quality of optical flow, scene flow, and stereo models advanced for over three decades,  
 102 their robustness recently regained attention as result of brittle deep learning generalization (Ranjan  
 103 et al., 2019; Schmalfuss et al., 2022b). We review robustness in dense-matching, particularly image  
 104 corruptions and metrics.

**105 Robustness in Dense Matching.** Robustness research for optical flow, scene flow, and stereo models  
 106 often focuses on *adversarial attacks*, which quantify prediction errors for optimized image pertur-  
 107

108 bations. Most attacks are for optical flow (Agnihotri et al., 2024c; Schmalfuss et al., 2023; 2022b; 109 Schrodi et al., 2022; Ranjan et al., 2019; Yamanaka et al., 2021; Koren et al., 2022) rather than 110 stereo (Berger et al., 2022; Wong et al., 2021) and scene flow (Wang et al., 2024a; Mahima et al., 111 2025). As remedies to adversarial vulnerability (Agnihotri et al., 2024b;a; 2023; Schrodi et al., 2022; 112 Anand et al., 2020) may be overcome through specialized optimization (Scheurer et al., 2024), another 113 line of robustness research considers non-adversarial data shifts. Those come in two flavors: 114 *generalization across datasets*, i.e. the Robust Vision Challenge (<http://www.robustvision.net/>), and 115 *robustness to image corruptions*. Work on dense matching models typically reports generalization 116 (Mehl et al., 2023a; Teed & Deng, 2020; 2021; Lipson et al., 2021; Huang et al., 2022; Xu et al., 117 2022b) to several datasets, which span synthetic (Mehl et al., 2023b; Butler et al., 2012; Richter et al., 118 2017; Mayer et al., 2016; Dosovitskiy et al., 2015; Gaidon et al., 2016; Ranjan et al., 2020; Li et al., 119 2024) and real-world data (Geiger et al., 2012; Menze & Geiger, 2015; Kondermann et al., 2016; 120 Scharstein et al., 2014; Schöps et al., 2017), often in automotive contexts. While some datasets 121 contain image corruptions, e.g. motion blur, depth of field, fog, noise, or brightness changes (Sun et al., 122 2021; Butler et al., 2012; Mehl et al., 2023b; Menze & Geiger, 2015), they do not systematically 123 assess corruption robustness. Yet, in the wild, robustness to image corruptions is crucial. For optical 124 flow, systematic low light (Zheng et al., 2020) and weather datasets (Schmalfuss et al., 2022a; 125 2023) exist, and Schrodi et al. (2022); Yi et al. (2024) apply 2D image corruptions (Hendrycks & 126 Dietterich, 2019) to optical flow data. Beyond these isolated works on optical flow, no systematic 127 image-corruption study before RobustSpring spans all three dense matching tasks and includes scene 128 flow or stereo.

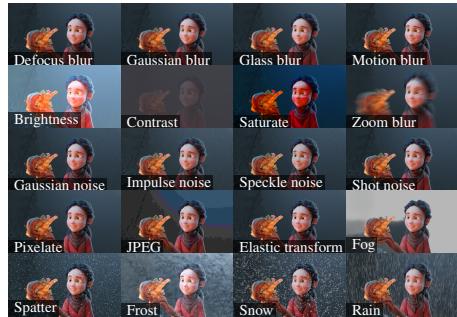
129 **Robustness to Image Corruptions.** Popularized by 2D common corruptions (Hendrycks & 130 Dietterich, 2019), the field of image corruption robustness rapidly expanded from classification 131 (Hendrycks & Dietterich, 2019; Müller et al., 2023) to depth estimation (Kar et al., 2022), 3D 132 object detection (Michaelis et al., 2019; Kong et al., 2023), and semantic segmentation (Kong et al., 133 2023). Conceptually, corruptions were extended to the 3D space (Kar et al., 2022), LiDAR (Kong 134 et al., 2023), and procedural rendering (Drenkow & Unberath, 2024), but none have been tailored to 135 the depth-, stereo-, and time-dependent setup of dense matching with optical flow, scene flow, and stereo.

136 **Robustness Metrics and Benchmarks.** Most robustness metrics for dense matching differ by 137 whether they utilize ground truth (Ranjan et al., 2019; Agnihotri et al., 2024c; Yi et al., 2024) or 138 not (Schmalfuss et al., 2022b; 2023; 2022a). However, multiple works (Schmalfuss et al., 2022b; 139 Tsipras et al., 2019; Taori et al., 2020) show that robustness and accuracy are competing qualities 140 that should not be quantified together. This informs our robustness metric. RobustSpring is the first 141 dense-matching *robustness* benchmark, and joins prior classification robustness benchmarks (Croce 142 et al., 2021; Jung et al., 2023; Tang et al., 2021)

### 3 ROBUSTSPRING DATASET AND BENCHMARK

147 RobustSpring is a large, novel image corruption dataset for optical flow, scene flow, and stereo. 148 Below, we describe how we build on Spring’s stereo video dataset and augment its frames with 149 diverse image corruptions integrated in time, stereo, and depth, how we evaluate robustness to image 150 corruptions, and use it to benchmark algorithm capabilities.

151 **Spring Data.** Spring (Mehl et al., 2023b) is a high-resolution benchmark with rendered stereo 152 sequences and dense ground truth. It is the ideal base for an image corruption dataset as its detailed 153 renderings permit image alterations of varying granularity – from removing detail by blurring to 154 adding detail via weather. Spring provides a public training and closed test split, where test ground 155 truth for optical flow, disparity, and extrinsic camera parameters is withheld. As RobustSpring 156 is designed to complement accuracy analyses, we build on the 2000 Spring test frames (two per 157 stereo camera). To apply corruptions with time, stereo, and depth consistency, we require depth 158 and extrinsics that are not publicly available. We therefore estimate extrinsics using COLMAP 3.8 159 and depths via  $Z = \frac{f_x \cdot B}{d}$ , with focal length  $f_x$ , baseline  $B$ , and disparities  $d$  predicted by MS- 160 RAFT+ (Jahedi et al., 2022; 2024). This estimation avoids ground-truth leakage while maintaining 161 benchmark integrity. Quantitative results on the accuracy of our depth and extrinsics estimation are 162 given in App. A.5, and a detailed discussion of motion ranges in Spring is provided in App. A.7.



(a) Image corruptions on a single image.

	Color	Blur	Noise	Qual	Weather
Property	Brightness Contrast Saturate	Defocus Gaussian Glass Motion Zoom	Gaussian Impulse Speckle Shot	Pixelate JPEG Elastic	Spatter Frost Snow Rain Fog
Time-cons.	✓✓✓	✓✓✓✓✓✓✓	---	✓✓✓	✓✓✓✓✓✓✓
Stereo-cons.	✓✓✓	✓✓✓	---	✓✓✓	---
Depth-cons.	---	---	---	---	---
SSIM	0.70 0.70	0.70 0.73	0.20 0.20	0.70 0.70	0.72 0.73 0.70 0.70

(b) Overview of corruptions and their consistency in time, stereo, or depth, with resulting visual changes w.r.t. the original images as SSIM.

Figure 2: Overview of RobustSpring’s image corruptions.

### 3.1 CORRUPTION DATASET CREATION

RobustSpring corrupts the Spring test frames via 20 diverse image corruptions, summarized in Fig. 2a and Fig. 2b. Below, we describe the image corruption types, their new consistencies, their implementation, and their severity levels.

**Corruption Types.** In RobustSpring, we consider the five image corruption types from Hendrycks & Dietterich (2019): color, blur, noise, quality, and weather. Color simulates different lighting conditions and camera settings, including brightness, contrast, and saturation. Blur acts like focus and motion artifacts, including defocus, Gaussian, glass, motion, and zoom blur. Noise represents sensor errors and ambience, including Gaussian, impulse, speckle, and shot noise. Quality distortions are lossy compressions and geometric distortions, including pixelation, JPEG, and elastic transformations. Weather enacts outdoor conditions, including spatter, frost, snow, rain, and fog. All corruptions applied to the same frame are shown in Fig. 2a.

While these 20 corruptions do not cover the entire corruption space, they are chosen to represent the most common perturbations encountered in real imagery and to provide a balanced basis for robustness evaluation. Several implementations were adapted and we additionally include practically relevant corruptions such as saturation, Gaussian blur, speckle noise, rain, and spatter. The blur and noise families span the most dominant optical and sensor degradations, and the weather and quality corruptions capture major outdoor and compression-related effects. At the same time, some perturbations, such as illumination changes requiring re-rendering (*e.g.* colored or dynamic lighting), lie beyond what can be approximated in post-processing. These, along with further outdoor effects (*e.g.* bloom, glare, dusty conditions) or extended codec distortions (*e.g.* JPEG 2000), form natural directions for future extensions.

**Corruption Consistencies.** To increase the realism of these 20 corruptions for dense matching models, we extend their definition to time, stereo, and depth: *Time-consistent* corruptions evolve smoothly over subsequent frames for a single camera, mirroring persistent lens or sensor effect, *e.g.* frost follows a temporally coherent pattern on one camera but differs between left and right. *Stereo-consistent* corruptions equally influence both stereo cameras, such as shared brightness or contrast adjustments. Unlike simply using the same hyperparameters, stereo consistency does not imply identical pixel-wise noise realizations, only that both views undergo the same transformation strength. *Depth-consistent* corruptions are rendered directly in the 3D scene, ensuring that their projection into each stereo view respects geometry. This applies to weather effects, such as snow, rain, and fog, where particles follow 3D trajectories and generate view-dependent projections. Other corruptions, such as blur or noise, do not benefit from 3D rendering; therefore, they use independent realizations per frame, despite sharing global severity parameters. Fig. 2b summarizes the consistencies we added to 16 of our 20 corruptions. Note that motion blur is not stereo-consistent because it depends on the specific camera view.

**Corruption Implementation.** Though most corruptions are loosely based on Hendrycks & Dietterich (2019), our corruption consistencies require multiple adaptations. Furthermore, we employ specialized techniques for highly consistent effects, *i.e.* motion blur, elastic transform, snow, rain and fog. We adapt implementations from Hendrycks & Dietterich (2019), modify glass blur, zoom blur, frost, and pixelation to accommodate higher resolutions and non-square images, and adjust

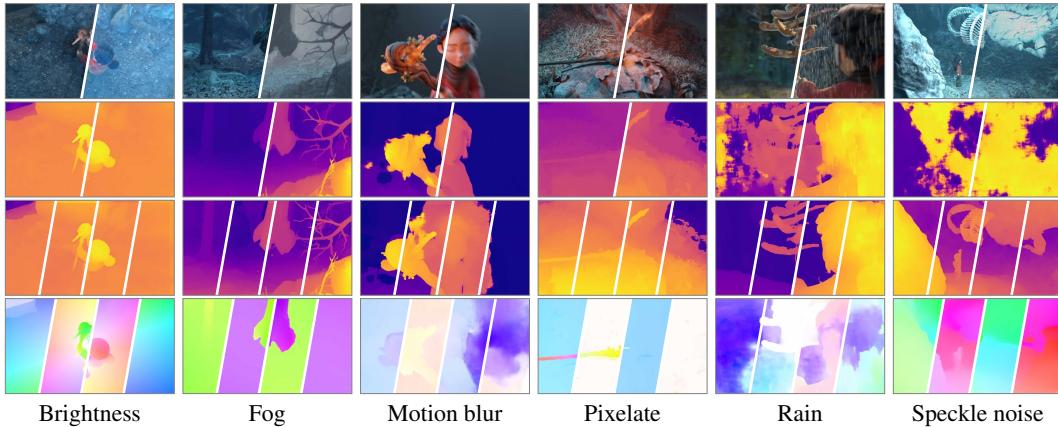


Figure 3: RobustSpring example frames. First row shows clean and corrupted images. Second row shows the left and right disparity maps predicted with LEA Stereo (Cheng et al., 2020). Third row shows the target disparities for forward left, backward left, forward right, and backward right directions from M-FUSE (Mehl et al., 2023a). Fourth row shows optical flow estimates for forward left, backward left, forward right, and backward right from RAFT (Teed & Deng, 2020). All disparities and flows are computed on the corrupted dataset. See Fig. 6 in App. A.3.1 for additional frames.

frost, glass blur, and spatter for consistency across video scenes. Motion blur is based on Zheng et al. (2006) and adds camera-induced motion with clean optical flow estimates. Elastic transform uses PyTorch’s transforms package to create a see-through-water-like effect, changing object morphology with smooth frame transitions. For snow and rain, we expand the two-step 3D particle rendering of Schmalfuss et al. (2023) to multi-step particle trajectories and stereo views, change from additive-blending to order-independent alpha blending (McGuire & Bavoil, 2013), and include global illumination (Halder et al., 2019). To augment the large-scale Spring data, we improve its performance via more effective particle generation and parallel processing. Fog is based on the Koschmieder model following Wiesemann & Jiang (2016). Full details are in App. A.3.3.

**Corruption Severity.** Prior works (Hendrycks & Dietterich, 2019; Müller et al., 2023; Kar et al., 2022; Michaelis et al., 2019; Kong et al., 2023) defined corruptions with several levels of severity. Here, we opt for one severity per corruption, because evaluating one scene flow model on all 20 corruptions already produces 2.1 TB of raw data – 1.2 GB after subsampling, *cf.* Sec. 3.2. More severity levels would overburden the evaluation resources of benchmark users. To balance severity across corruptions, we tune their hyperparameters until the image SSIM reaches a defined threshold. We generally use  $\text{SSIM} \geq 0.7$ , and, because the SSIM is less sensitive to blurs than noises (Hore & Ziou, 2010),  $\text{SSIM} \geq 0.2$  for noises for visually similar artifact strengths. We conducted a focused perceptual study to validate the selection of corruption strengths and their corresponding SSIM values in RobustSpring. See App. A.6 for more details. Final SSIMs are in Fig. 2b.

### 3.2 ROBUSTNESS EVALUATION METRIC

With various corruption types, we need a metric to quantify model robustness to these variations. In the following, we motivate and derive a ground-truth-free robustness metric for dense matching, introduce subsampling for efficiency, and discuss strategies for joint rankings over corruptions.

**Definition of Optical Flow.** Throughout this work, and consistent with the Spring benchmark (Mehl et al., 2023b), we define optical flow as the *true 3D motion of visible surfaces projected into the 2D image plane* (Horn & Schunck, 1981; Baker et al., 2011). This definition is standard in classical benchmarks and provides the basis for accuracy evaluation. Alternative formulations, such as apparent motion, exist in the literature, but RobustSpring’s robustness evaluation is agnostic to this choice, since robustness is measured independently of ground truth. A more detailed discussion on the implications of this definition, including the relation between perfect accuracy and perfect robustness, is provided in App. A.1.

**Robustness Metric Concepts.** For dense matching, robustness to corruptions lacks a standardized evaluation metric. Metrics exist for adversarial robustness, using the distance between corrupt

prediction and either (i) ground-truth (Ranjan et al., 2019; Agnihotri et al., 2024c) or (ii) clean prediction (Schmalfuss et al., 2022b; 2023; 2022a). The latter is preferred for two reasons: First, (i)’s ground-truth comparisons mix accuracy and robustness, which are competing model qualities (Schmalfuss et al., 2022b; Tsipras et al., 2019; Taori et al., 2020) that should be separate. This competition is intuitive: A model that always outputs the same value is as robust as inaccurate. Likewise, an accurate model varies for any input change and thus is not robust. Second, (ii) separates robustness from accuracy and builds on an established mathematical concept for system robustness (Hein & Andriushchenko, 2017; Pauli et al., 2022): the Lipschitz constant  $L^c$ . It defines robust models as those whose prediction  $f$  is similar on clean and corrupt image  $I$  and  $I^c$ , relative to their difference. For dense matching, it reads

$$L^c = \frac{\|f(I) - f(I^c)\|}{\|I - I^c\|}, \quad (1)$$

where the term  $\|I - I^c\|$  refers to the per-pixel intensity difference between the clean and corrupted images. This robustness formulation is preferable for real-world applications that demand stable scene estimations *despite* corruptions like snow, and remains valid independent of whether optical flow is defined as true motion or apparent motion. We emphasize that RobustSpring explicitly measures robustness in terms of *stability*: models are considered robust if their predictions remain consistent under corrupted inputs. Lower robustness scores correspond to higher stability, not improved accuracy. Other definitions of robustness, such as ground-truth-based robustness, are possible and may be more appropriate for different use cases. Our benchmark deliberately focuses on stability-based robustness, providing a complementary axis of evaluation alongside accuracy.

**Corruption Robustness Metric.** Based on Eq. (1), we quantify model robustness to corruptions. Because RobustSpring’s corrupt images  $I^c$  deviate from their clean counterparts  $I$  by a similar amount, *cf.* SSIM equalization in Sec. 3.1, we omit the denominator in Eq. (1) and define *corruption robustness*  $R^c$  as distance between clean  $f(I)$  and corrupted  $f(I^c)$  predictions with distance metric  $M$ :

$$R_M^c = M[f(I), f(I^c)]. \quad (2)$$

For similarity to Spring’s evaluation, we use corruption robustness with various metrics  $M$ , reporting  $R_{\text{EPE}}^c$ ,  $R_{\text{1px}}^c$  and  $R_{\text{Fl}}^c$  for optical and scene flow, and  $R_{\text{1px}}^c$ ,  $R_{\text{Abs}}^c$  and  $R_{\text{D1}}^c$  for stereo. Here, EPE denotes the average end-point error, 1px the percentage of pixels with an error exceeding 1 px, Fl the KITTI optical-flow outlier rate, Abs the mean absolute disparity error, and D1 the KITTI disparity outlier rate, see Mehl et al. (2023b) and Menze & Geiger (2015) for more details. Interestingly, our EPE-based corruption robustness

$$R_{\text{EPE}}^c = \text{EPE}[f(I), f(I^c)] = \frac{1}{|\Omega|} \sum_{i \in \Omega} \|f_i(I) - f_i(I^c)\| \quad (3)$$

on image domain  $\Omega$  is a generalization of optical-flow adversarial robustness (Schmalfuss et al., 2022b) to dense matching and corruptions.

**Metric Subsampling.** For a benchmark, users should upload robustness results to a web server. Given the large number of 20 datasets, data reduction is essential to facilitate evaluations and uploads. To this end, we evaluate on a reduced set of pixels by refining the original subsampling strategy from Spring, which retains about 1% of the full data. First, we additionally subsample the set of full-resolution Hero-frames because they are only kept in the original benchmark for visualization purposes, leaving 0.95%, and then apply 20-fold subsampling, ultimately keeping 0.05% of the full data.

**Robustness Ranking.** Because we generate 20 different corruption evaluations *per* dense matching model, we need a summarization strategy to produce one result per model. Per-model results are ranked based on three strategies: Average, Median, and the Schulze voting method (Schulze, 2018). In contrast to averaging across all 20 evaluations, the median reduces the impact of extreme outliers. The Schulze method provides a holistic, pairwise comparison approach that ranks models based on preference aggregation and was used for prior generalization evaluations in the Robust Vision Challenges. We evaluate their differences in Sec. 4.2.

### 3.3 DATASET AND BENCHMARK FUNCTIONALITY

Below, we summarize RobustSpring’s corruption dataset and describe its benchmark function. Fig. 3 shows data samples with stereo, optical flow, and scene flow estimates.

324 **Table 1:** Initial RobustSpring results on corruption robustness of optical flow models, using  $R_{\text{EPE}}^c$ ,  
 325  $R_{\text{Ipx}}^c$  and  $R_{\text{Fl}}^c$  between clean and corrupted flow predictions. Low values indicate robust models.  
 326  $\varepsilon_{\text{clean}}$  compares clean predictions with ground-truth flow, values from Mehl et al. (2023b).

	SEA-RAFT			GMFlow			MS-RAFT+			FlowFormer			GMA			SPyNet			RAFT			FlowNet2			PWCNet			
	$R_{\text{EPE}}^c$	$R_{\text{Ipx}}^c$	$R_{\text{Fl}}^c$	$R_{\text{EPE}}^c$	$R_{\text{Ipx}}^c$	$R_{\text{Fl}}^c$	$R_{\text{EPE}}^c$	$R_{\text{Ipx}}^c$	$R_{\text{Fl}}^c$	$R_{\text{EPE}}^c$	$R_{\text{Ipx}}^c$	$R_{\text{Fl}}^c$	$R_{\text{EPE}}^c$	$R_{\text{Ipx}}^c$	$R_{\text{Fl}}^c$	$R_{\text{EPE}}^c$	$R_{\text{Ipx}}^c$	$R_{\text{Fl}}^c$	$R_{\text{EPE}}^c$	$R_{\text{Ipx}}^c$	$R_{\text{Fl}}^c$	$R_{\text{EPE}}^c$	$R_{\text{Ipx}}^c$	$R_{\text{Fl}}^c$				
Color	Brightness	0.21	1.65	0.40	0.33	3.31	1.12	0.33	2.88	1.02	0.68	2.82	1.05	0.36	3.22	1.04	2.72	14.67	8.91	0.92	3.49	1.61	0.45	3.16	1.05	1.04	7.38	3.00
	Contrast	0.75	3.75	1.51	0.46	6.71	1.71	0.87	6.69	3.24	0.93	5.48	1.96	0.68	6.43	2.20	8.23	38.90	27.23	1.32	5.73	2.64	1.87	9.26	4.74	2.98	30.07	7.42
	Saturation	0.16	1.29	0.36	0.34	3.30	0.96	0.34	2.87	1.03	0.42	2.39	0.88	0.43	3.47	1.18	3.36	17.34	11.31	0.93	3.33	1.47	0.51	3.40	1.10	1.21	9.92	3.68
	Blur	Defocus	0.20	1.47	0.42	0.53	6.17	1.45	0.51	4.01	1.47	0.55	3.85	1.19	0.56	5.02	2.01	0.57	10.16	1.36	1.03	4.70	2.07	0.53	3.35	1.06	0.98	6.51
Blur	Gaussian	0.23	1.79	0.51	0.66	7.77	1.88	0.58	4.45	1.63	0.63	4.32	1.37	0.62	5.48	2.22	0.76	15.44	2.12	1.10	5.12	2.26	0.60	4.05	1.27	1.11	7.72	3.09
	Glass	0.22	1.46	0.41	0.85	20.87	1.82	0.53	4.45	1.37	0.64	4.04	1.17	0.61	5.60	1.91	0.75	16.94	1.36	1.08	5.13	1.97	0.50	3.12	0.96	0.91	5.96	2.47
	Motion	1.11	13.78	5.82	1.38	18.35	7.51	1.31	14.06	6.16	1.35	14.03	5.77	1.19	14.40	6.18	2.32	19.55	10.05	2.08	14.33	6.35	1.60	14.07	6.47	1.95	16.25	7.47
	Zoom	1.83	24.51	8.41	1.88	35.80	9.90	1.81	21.84	7.13	1.66	22.72	6.77	1.54	23.17	7.16	4.82	46.67	28.37	3.14	22.80	7.61	2.36	24.61	9.04	3.52	50.33	15.64
Noise	Gaussian	1.49	12.30	4.34	4.70	57.95	21.67	5.70	35.74	22.12	5.66	27.83	18.30	2.81	24.70	12.96	2.22	42.23	14.88	7.43	27.92	18.99	1.33	11.24	5.06	2.79	26.87	9.89
	Impulse	2.58	19.01	7.51	6.64	66.14	28.70	7.39	45.72	29.05	7.33	23.58	14.47	4.08	31.31	18.13	2.92	53.45	20.41	6.51	29.65	18.32	2.37	15.70	7.48	3.57	35.67	14.45
	Shot	1.29	12.15	3.61	3.90	62.01	20.64	4.22	34.98	17.18	5.47	25.52	15.60	5.32	22.52	12.66	1.95	46.32	12.89	6.62	26.05	16.48	1.32	12.57	4.19	2.74	26.83	8.00
	JPEG	1.20	10.69	3.41	3.52	56.71	17.77	4.36	31.67	17.77	5.75	26.02	16.01	3.15	23.11	11.59	7.64	25.64	17.08	1.16	9.87	3.92	2.59	23.75	7.88			
Pixelation	Pixelated	0.41	1.88	0.51	19.6	68.09	18.71	1.60	45.83	6.78	1.48	31.68	2.59	1.11	25.86	1.78	1.22	50.63	2.90	1.65	21.47	2.00	0.77	7.74	0.88	0.92	8.67	2.22
	JPEG	4.02	34.86	13.57	3.32	83.54	27.92	2.09	41.61	12.82	2.82	42.42	14.96	1.92	38.70	11.51	2.95	53.97	18.08	3.19	37.72	13.67	2.56	31.09	11.85	2.88	49.15	15.91
	Elastic	0.56	11.14	1.42	1.37	40.00	6.89	1.16	32.49	5.54	2.62	35.78	11.01	1.24	27.24	6.40	1.08	34.62	4.77	1.33	19.43	4.78	0.79	16.27	2.12	1.42	28.18	5.47
	Fog	1.20	11.20	7.41	0.80	14.42	5.32	0.91	10.32	6.33	0.86	9.66	5.67	0.84	11.21	6.42	5.20	28.15	19.97	1.97	12.01	7.11	1.74	11.77	7.82	16.84	20.96	12.89
Weather	Frost	7.60	36.86	21.86	8.20	63.96	29.96	7.38	29.96	21.25	8.18	34.19	23.87	8.13	34.20	22.31	6.97	45.13	10.13	8.37	32.75	21.76	7.22	33.69	21.15	8.27	50.31	27.44
	Rain	16.73	36.87	25.09	8.60	64.20	32.72	19.99	36.74	31.22	11.13	33.50	20.83	33.00	43.98	36.18	18.20	68.87	56.65	42.41	38.89	31.99	63.71	48.25	41.15	40.18	73.51	57.05
	Snow	8.54	60.52	43.81	3.60	70.60	29.90	4.69	33.21	30.91	7.92	40.20	33.87	12.08	74.27	66.65	7.16	37.04	31.37	39.79	68.67	61.60	39.73	90.80	81.91			
	Spatter	8.93	53.31	30.52	6.58	67.90	27.09	6.63	28.22	20.24	8.41	40.38	26.92	7.75	36.11	21.81	5.71	48.60	33.82	7.99	30.37	19.87	9.13	45.03	28.99	9.33	65.41	40.19
Average	<b>2.96</b>	<b>17.52</b>	<b>9.05</b>	<b>2.98</b>	<b>40.89</b>	<b>14.68</b>	<b>3.62</b>	<b>23.29</b>	<b>12.21</b>	<b>3.77</b>	<b>21.53</b>	<b>11.21</b>	<b>4.03</b>	<b>21.47</b>	<b>10.95</b>	<b>4.29</b>	<b>38.32</b>	<b>19.18</b>	<b>5.64</b>	<b>20.18</b>	<b>11.47</b>	<b>7.01</b>	<b>18.84</b>	<b>11.09</b>	<b>7.25</b>	<b>31.71</b>	<b>16.44</b>	
	Std. Dev.	4.29	17.98	12.08	2.70	27.91	11.91	4.58	15.54	10.62	3.44	14.37	9.94	7.23	13.67	10.55	4.38	18.35	17.60	9.10	12.55	9.98	15.94	17.93	15.87	11.83	24.43	20.79
Median	<b>1.20</b>	<b>11.68</b>	<b>3.98</b>	<b>1.92</b>	<b>48.35</b>	<b>13.83</b>	<b>1.71</b>	<b>29.09</b>	<b>6.95</b>	<b>2.14</b>	<b>24.55</b>	<b>8.89</b>	<b>1.39</b>	<b>23.93</b>	<b>6.79</b>	<b>2.82</b>	<b>41.33</b>	<b>13.88</b>	<b>2.60</b>	<b>22.13</b>	<b>7.36</b>	<b>1.47</b>	<b>12.17</b>	<b>4.90</b>	<b>2.77</b>	<b>26.85</b>	<b>7.94</b>	
	$\varepsilon_{\text{clean}}$ (Clean Error)	0.36	3.69	1.35	0.94	10.36	2.95	0.64	5.72	2.19	0.72	6.51	2.38	0.91	7.07	3.08	4.16	29.96	12.87	1.48	6.79	3.20	1.04	6.71	2.82	2.29	82.27	4.89

342 **RobustSpring Dataset.** The final RobustSpring dataset entails 20 corrupted versions of Spring,  
 343 resulting in 40,000 frames, or 20,000 stereo frame pairs. Each corruption evaluation yields 3960  
 344 optical flows (990 per camera & direction), 2000 stereo disparities (1000 per camera), and 3960  
 345 additional scene flow disparity maps (990 per camera per direction). To discourage corruption  
 346 finetuning and ensure a fair benchmark, we do not provide corrupted training data, only corrupted test  
 347 data. We separately provide the raw data and a curated dataset for predicting dense matches.

348 **RobustSpring Benchmark.** RobustSpring enables uploading robustness results to a benchmark  
 349 website for display in a public ranking. To emphasize that robustness and accuracy are two axes  
 350 of model performance with equal importance (Tsipras et al., 2019), we couple RobustSpring with  
 351 Spring’s established accuracy benchmark. Thus, researchers can report model robustness and accu-  
 352 racy on the same dataset. We provide mock-ups of this integration in App. A.8.

## 4 RESULTS

356 We evaluate RobustSpring under two aspects: First, we report initial results for 17 optical flow,  
 357 scene flow and stereo models. Then, we analyze the benchmark evaluation, particularly subsampling  
 358 strategy and ranking methods.

### 4.1 INITIAL ROBUSTSPRING BENCHMARK RESULTS

362 We provide initial results on RobustSpring for selected models from all three dense matching  
 363 tasks. For optical flow, we include SEA-RAFT (Wang et al., 2024b), GMFlow (Xu et al., 2022b),  
 364 MS-RAFT+ (Jahedi et al., 2024), FlowFormer (Huang et al., 2022), GMA (Jiang et al., 2021),  
 365 SPyNet (Ranjan & Black, 2017), RAFT (Teed & Deng, 2020), FlowNet2 (Ilg et al., 2017), and  
 366 PWCNet (Sun et al., 2018). For scene flow, we evaluate M-FUSE (Mehl et al., 2023a) and RAFT-  
 367 3D (Teed & Deng, 2021). For stereo estimation, we evaluate RAFT-Stereo (Lipson et al., 2021),  
 368 ACVNet (Xu et al., 2022a), LEAStereo (Cheng et al., 2020), and GANet (Zhang et al., 2019). An  
 369 overview of all models and used checkpoints is in the Appendix in Tab. 4. Importantly, none of these  
 370 models are fine-tuned to either Spring or RobustSpring data, to assess the generalization capacity of  
 371 existing models.

372 **Optical Flow.** The evaluation results in Tab. 1 and Fig. 4a show large robustness variations across  
 373 corruption types. Weather corruptions, especially rain and snow, degrade performance most, while  
 374 color corruptions have little effect. Model rankings also vary: FlowNet2 performs poorly overall but  
 375 is the most resilient to noise (Fig. 4b). SEA-RAFT and GMFlow achieve the lowest average  $R_{\text{EPE}}^c$ ,  
 376 while GMA yields the lowest median, as detailed in Sec. 4.2.

377 To examine accuracy–robustness relations, we compare both in Fig. 4c. Accurate models tend to be  
 378 more robust, though no method excels in both dimensions, and the trend is weak and nonlinear (see

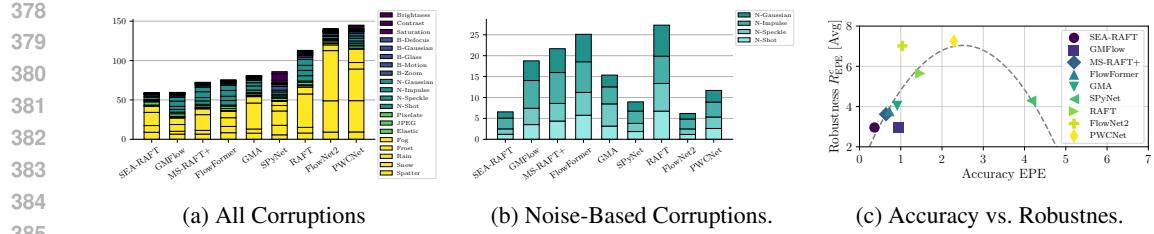


Figure 4: Accumulated corruption robustness  $R_{\text{EPE}}^c$  for optical flow models over all corruptions [left], only noise corruptions [middle], and accuracy vs. robustness [right] with a dashed line representing a quadratic polynomial fit. Small values are robust (and accurate) models. All other corruption classes color (purple), blur (blue), noise (cyan), quality (green), and weather (yellow) are in App. A.4 in Fig. 7, and Fig. 8 shows accuracy vs. Median  $R_{\text{EPE}}^c$ .

also median robustness in Fig. 8). Fitting a quadratic regression to Fig. 4c highlights this nonlinearity. Accuracy and robustness coincide for some corruption families but diverge substantially for others. This stands in contrast to adversarial robustness, where a strong accuracy–robustness trade-off is observed (Schmalfuss et al., 2022b). The key to understanding these differences is to consider different corruption types individually. Because accurate models need to process fine details in order to achieve highly accurate predictions, this attention to details can be exploited by noise-based adversarial attacks in Schmalfuss et al. (2022b). For RobustSpring’s non-adversarial corruptions, accurate models are also relatively sensitive to noise corruptions (cf. Fig. 4b). However, the average corruption robustness is mostly dominated by weather corruptions (cf. Fig. 4a), where more accurate models achieve better robustness values. We identify the reason for this in Sec. 4.2: accurate models retain corruption-induced motion errors around the particles rather than spreading them onto the background (cf. Fig. 5a). This aligns with the results in Schmalfuss et al. (2023), which found accurate models to be more robust towards adversarial weather attacks.

Architectural trends also emerge. Transformer models (GMFlow, FlowFormer) perform best overall but struggle with noise, likely due to global matching. Hierarchical models (e.g. MS-RAFT+) show balanced robustness and may benefit from multi-scale feature processing to cope with quality degradations. Stacked models (e.g. SEA-RAFT, FlowNet2) uniquely resist noise, possibly due to their progressive refinement across layers. Overall, architecture influences robustness to specific corruptions, but no paradigm is universally superior.

**Scene Flow.** The results for scene flow are in Tab. 2a, and include optical flow and target frame disparity predictions for M-FUSE and RAFT-3D. M-FUSE generally produces more robust optical flow across corruptions with a lower average  $R_{\text{EPE}}^c$  than RAFT-3D. But both methods suffer sig-

Table 2: Initial RobustSpring results on corruption robustness of scene flow and stereo disparity models, using corruption robustness  $R_{1\text{px}}^c$ ,  $R_{\text{Abs}}^c$  and  $R_{\text{DI}}^c$  between clean and corrupted predictions. Low values indicate robust models. Corresponding Disparity 1 from scene flow models LEAStereo (s) for M-FUSE, and GANet (s) for RAFT-3D in Tab. 2b. Stereo disparity models use Stereo (s) and KITTI (k) checkpoints.

Color	M-FUSE										RAFT-3D																			
	Optical flow					Disparity 2					Optical flow					Disparity 2														
	$R_{\text{EPE}}^c$	$R_{1\text{px}}^c$	$R_{\text{Abs}}^c$	$R_{\text{DI}}^c$																										
Brightness	0.83	5.54	2.90	0.14	0.53	0.18	0.38	0.07	1.48	0.21	0.99	7.86	3.60	0.17	0.17	1.42	10.71	5.07	0.07	1.65	0.22									
Contrast	0.67	4.94	2.43	0.12	1.22	0.14	0.93	0.67	3.31	0.03	1.33	0.18	5.26	2.80	0.17	0.17	1.22	10.71	5.07	0.06	1.34	0.18								
Defocus	0.84	5.26	2.71	0.15	1.37	0.15	0.66	5.27	2.44	0.06	0.88	0.10	5.98	3.81	0.16	0.16	1.56	0.18	0.58	2.73	0.06	1.04	0.14							
Gaussian	0.98	5.81	2.92	0.16	1.56	0.18	0.78	5.20	2.87	0.06	0.98	0.10	6.91	3.77	0.17	0.17	1.77	10.71	5.07	0.04	1.09	0.10								
Blur	0.98	5.17	2.80	0.16	1.56	0.18	0.78	5.20	2.87	0.06	0.98	0.10	1.51	1.10	0.81	0.18	2.50	0.35	1.62	14.66	6.85	0.08	1.60	0.28						
Motion	4.87	26.32	12.34	0.36	5.60	0.69	4.65	42.07	22.91	0.18	3.26	0.67	2.28	7.92	0.28	3.74	34.06	11.99	0.14	2.84	0.28									
Zoom	2.28	27.88	9.52	0.28	3.74	0.28	2.68	34.06	11.99	0.14	2.84	0.28	0.98	1.76	0.22	0.22	2.28	11.99	0.14	0.22	0.22	0.22								
Gaussian	6.49	29.22	14.81	0.41	6.56	0.80	5.25	43.33	25.43	0.20	3.64	0.71	5.98	37.32	19.16	0.43	5.25	33.16	19.00	0.33	3.64	0.71								
Impulse	5.98	37.32	19.16	0.43	8.11	0.86	6.73	59.86	33.16	0.22	4.44	0.75	6.73	59.86	33.16	0.43	8.11	59.86	33.16	0.22	4.44	0.75								
Speckle	3.73	29.39	12.22	0.35	6.56	0.57	4.86	51.12	26.11	0.16	3.17	0.64	1.51	10.10	6.81	0.18	2.50	0.35	1.62	14.66	6.85	0.08	1.60	0.28						
Shot	4.87	26.32	12.34	0.36	5.60	0.69	4.65	42.07	22.91	0.18	3.26	0.67	2.28	7.92	0.28	3.74	34.06	11.99	0.14	2.84	0.28									
Pixelated	0.86	5.95	2.51	0.19	1.51	0.13	0.82	7.66	2.83	0.06	1.02	0.10	0.98	6.72	3.60	0.20	3.40	33.33	19.55	0.20	3.40	0.20								
JPEG	0.98	5.17	2.80	0.16	3.40	0.36	0.82	33.33	19.55	0.20	3.40	0.20	0.98	6.72	3.60	0.20	3.40	33.33	19.55	0.20	3.40	0.20								
Elastic	1.15	23.93	9.02	0.22	2.28	0.22	1.76	18.2	7.04	0.08	1.61	0.20	0.98	6.72	3.60	0.20	3.40	33.33	19.55	0.20	3.40	0.20								
Weather	2.55	15.39	10.13	0.19	2.43	0.19	2.29	15.11	11.67	0.06	1.23	0.15	7.91	41.60	23.41	0.38	6.55	0.78	45.07	24.26	0.16	3.75	0.52							
Fog	7.91	41.60	23.41	0.38	6.55	0.78	7.49	45.07	24.26	0.16	3.75	0.52	10.21	41.78	28.99	0.70	12.79	74.23	59.77	0.47	10.75	1.96								
Rain	6.36	47.06	35.00	0.46	7.67	0.86	19.08	80.49	60.01	0.31	6.79	0.84	7.06	55.55	55.55	0.76	55.55	28.00	1.07	3.82	0.53									
Snow	7.00	46.35	22.10	0.39	6.21	0.86	7.06	55.55	25.80	0.01	3.82	0.53	1.15	23.93	9.02	0.22	1.76	18.2	7.04	0.08	1.61	0.20								
Spatter	<b>Average</b>	<b>3.39</b>	<b>11.67</b>	<b>11.17</b>	<b>0.29</b>	<b>4.20</b>	<b>0.46</b>	<b>5.03</b>	<b>32.16</b>	<b>17.36</b>	<b>0.14</b>	<b>2.89</b>	<b>0.46</b>	<b>36.80</b>	<b>15.79</b>	<b>16.81</b>	<b>40.95</b>	<b>20.90</b>	<b>20.77</b>	<b>50.02</b>	<b>8.24</b>	<b>14.04</b>	<b>44.71</b>	<b>6.44</b>	<b>10.86</b>	<b>45.26</b>	<b>12.11</b>	<b>17.93</b>		
Std. Dev.	2.93	15.23	9.60	0.34	3.85	0.34	1.85	32.16	17.36	0.11	2.40	0.43	20.16	20.72	10.79	23.56	18.64	17.09	29.19	31.33	23.64	23.66	21.37	17.25						
Median	<b>2.13</b>	<b>20.86</b>	<b>8.17</b>	<b>0.25</b>	<b>3.06</b>	<b>0.35</b>	<b>2.49</b>	<b>27.88</b>	<b>11.11</b>	<b>0.10</b>	<b>2.12</b>	<b>0.35</b>	<b>4.03</b>	<b>3.78</b>	<b>6.23</b>	<b>37.21</b>	<b>8.16</b>	<b>9.60</b>	<b>36.81</b>	<b>6.83</b>	<b>10.51</b>	<b>55.58</b>	<b>7.55</b>	<b>10.90</b>	<b>48.53</b>	<b>6.92</b>	<b>11.58</b>	<b>46.64</b>	<b>6.40</b>	<b>9.93</b>
Clean Error	2.52	13.96	6.89	7.11	32.95	14.54	2.53	20.98	8.48	8.08	57.03	21.54	15.27	3.02	5.35	14.77	1.52	5.35	19.89	3.88	9.19	47.50	6.15	17.16	27.91	5.29	11.56	23.22	4.59	10.39

432 nificant performance losses for severe weather like rain and noise-based corruptions, *e.g.* impulse  
 433 noise. Interestingly, their robustness does not improve compared to conventional optical flow mod-  
 434 els. Noise and weather corruptions remain a challenge for Disparity 2 predictions. Here, RAFT-3D  
 435 consistently achieves lower robustness scores compared to M-FUSE, but conditions like impulse  
 436 noise or rain still notably affect disparity predictions. Overall, both models have limited robustness,  
 437 but temporal consistency may contribute to lower robustness scores under several corruption types.

438 **Stereo.** Results of stereo disparity estimations are presented in Tab. 2b. The effect of the different  
 439 corruptions on the performance is significant, with noise and weather-based corruptions leading to  
 440 the largest errors, especially for GANet and LEAStereo. In particular, Gaussian and impulse noise  
 441 introduce extremely large errors, highlighting the sensitivity of stereo models to pixel-level noise.  
 442 Blur distortions, especially zoom blur, also have a severe impact on all models. In contrast, color-  
 443 based distortions generally yield smaller errors. RAFT-Stereo shows stronger resilience across most  
 444 corruption groups, performing better on color and noise based corruption than other models, but also  
 445 struggles with noise and severe weather effects such as rain and snow.

446 Table 3: Evaluations of the metrics used in RobustSpring.  
 447

448 (a) Influence of subsampling. We compare robustness  
 449 evaluations on the full test data (Full) to evaluations on  
 450 Spring’s original subsampling (Spring), original subsam-  
 451 pling without Hero-frames (Spring\*), and our refined cor-  
 452 ruption subsampling (Ours).

% Original Data	Subsampling $R_{\text{EPE}}^c$				Subsampling $R_{\text{Ipx}}^c$				
	Full	Spring	Spring*	Ours	Full	Spring	Spring*	Ours	
	100%	1.00%	0.94%	0.05%		100%	1.00%	0.94%	0.05%
SEA-RAFT	2.96	3.19	2.96	2.96	17.52	18.44	17.52	17.52	
GMFlow	2.98	3.20	2.98	2.98	40.89	41.99	40.89	40.89	
MS-RAFT+	3.62	3.84	3.62	3.62	23.38	24.44	23.39	23.39	
FlowFormer	3.77	3.89	3.77	3.77	21.52	22.39	21.53	21.53	
GMA	4.03	4.28	4.03	4.03	21.47	22.59	21.48	21.47	
SPyNet	4.30	4.56	4.29	4.29	38.32	39.28	38.32	38.32	
RAFT	5.64	6.15	5.64	5.64	20.17	21.20	20.18	20.18	
FlowNet2	7.01	7.36	7.01	7.01	18.84	19.79	18.84	18.84	
PWCNet	7.25	7.52	7.25	7.25	31.71	32.55	31.72	31.71	

448 (b) Robustness ranking of optical flow models  
 449 with ranking strategies Average  $R_{\text{EPE}}^c$ , Median  
 450  $R_{\text{EPE}}^c$ , and Schulze to summarize results over  
 451 corruptions. Please note that Schulze does not  
 452 produce numeric values.

Rank	Ranking Method		
	Average $R_{\text{EPE}}^c$	Median $R_{\text{EPE}}^c$	Schulze
1	2.96 SEA-RAFT	1.20 SEA-RAFT	SEA-RAFT
1	2.98 GMFlow	1.39 GMA	MS-RAFT+
2	3.62 MS-RAFT+	1.47 FlowNet2	GMA
3	3.77 FlowFormer	1.71 MS-RAFT+	FlowNet2
4	4.03 GMA	1.92 GMFlow	GMFlow
5	4.29 SPyNet	2.14 FlowFormer	FlowFormer
6	5.64 RAFT	2.60 RAFT	SPyNet
7	7.01 FlowNet2	2.77 PWCNet	PWCNet
8	7.25 PWCNet	2.82 SPyNet	RAFT

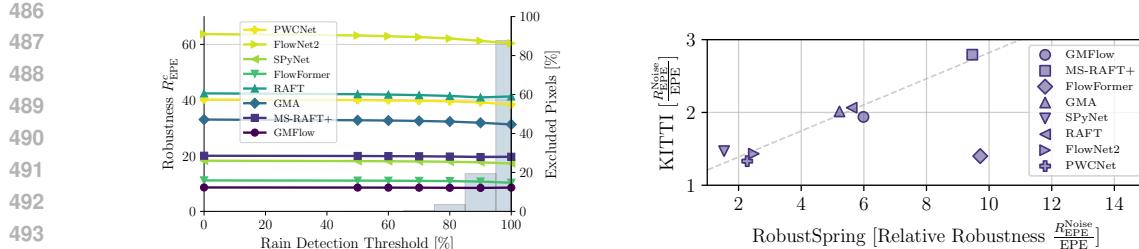
## 461 462 463 4.2 METRICS AND BENCHMARK CAPABILITY

464 After reporting initial RobustSpring results, we analyze aspects of its benchmark character: The  
 465 subsampling strategy for data efficiency, and different ranking systems for result comparisons across  
 466 20 different prompt variations. We also validate our robustness metric for object corruptions and  
 467 explore RobustSpring’s transferability to the real-world.

468 **Subsampling.** We evaluate RobustSpring’s strict data subsampling by comparing to results on the  
 469 full test set. As shown in Tab. 3a, our subsampling strategy produces results that are nearly identical  
 470 to those that include all pixels in the robustness calculation. We observe the largest discrepancy  
 471 for Spring’s original subsampling, because it includes a handful of full-resolution Hero-frames. If  
 472 those frames are also subsampled (Spring\*), results align with the full dataset. Overall, our stricter  
 473 subsampling to 0.05% of all data is not only data efficient but also exact.

474 **Metric Ranking.** To explore how ranking strategies influence the optical-flow robustness order,  
 475 we contrast our three summarization strategies: Average, Median, and Schulze (*cf.* App. A.4.3).  
 476 The rankings in Tab. 3b notably differ across strategies. The Average differs most from the other  
 477 rankings. For example, it ranks GMFlow 2nd, which is only 5th on Median and Schulze, suggesting  
 478 a good performance across corruptions without excessive outliers but no top performance on most  
 479 corruptions. Interestingly, Median and Schulze rankings are more aligned. As Schulze’s ranking  
 480 involves complex comparisons of per-corruption rankings and must be globally recomputed for new  
 481 models, the Median ranking is a cheap approximation to it. The ranking strategy has significant  
 482 implications for selecting robust models. While SEA-RAFT is optimal across rankings, the rankings  
 483 accentuate different aspects: overall performance, outlier robustness, or balanced performance in  
 484 pairwise comparisons. Hence, RobustSpring reports them all.

485 **Robustness on Object Corruptions.** In RobustSpring, robustness is defined as *stability* of pre-  
 486 dictions under corrupted inputs: models are robust if their outputs remain consistent for clean and



(a) Stability of corruption robustness  $R_{\text{EPE}}^c$  on rain corruption. Robustness scores and rankings remain stable even if no rain pixels are in the  $R_{\text{EPE}}^c$  calculation. (b) Relative robustness to noise on RobustSpring transfers to noisy real-world KITTI data (Menze & Geiger, 2015) for most optical flow models.

Figure 5: Additional evaluations of RobustSpring’s benchmark character.

corrupted scenes (Eq. 1). A natural question is whether this metric remains valid when corruptions introduce moving objects, such as rain or snow. To test this, we separate the contributions of background and corruption pixels by excluding pixels of objects like rain drops from the score calculation. Object pixels are detected via the value difference  $d$  between clean and corrupted images, and excluded if  $(1 - d)$  exceeds a threshold. Threshold 0 excludes none (the vanilla  $R_{\text{EPE}}^{\text{Rain}}$ ), while threshold 100 excludes all. Figure 5a shows the robustness score if rain is excluded from the calculation, along with bars indicating the amount [%] of excluded pixels. Remarkably, the robustness score changes by at most 5% even when all rain pixels (about 90% of the image) are discarded. High scores for rain or snow (*cf.* App. A.4.4) thus result mainly from mispredictions in the *periphery* of altered pixels, not from motion predictions on altered pixels. As scene-wide effects dominate, our stability-based robustness yields consistent rankings suited for broad robustness evaluations.

**Robustness in the Real World.** Finally, we investigate if RobustSpring’s corruption robustness transfers to the real world. To this end, we select the noisiest 10% KITTI data, estimating noise as in Immerkaer (1996). These noisy KITTI frames have no clean counterparts to calculate corruption robustness  $R_{\text{EPE}}^{\text{Noise}}$ . Thus, we approximate  $R_{\text{EPE}}^{\text{Noise}}$  via the accuracy difference on noisy and non-noisy KITTI frames. To account for model-specific performance differences on Spring and KITTI, we normalize with the clean dataset performance and show the resulting relative robustness  $\frac{R_{\text{EPE}}^{\text{Noise}}}{R_{\text{EPE}}^{\text{Clean}}}$  in Fig. 5b. Relatively robust models with low scores on RobustSpring are also robust on KITTI and vice versa. The only outlier, FlowFormer, overperforms on KITTI, potentially due to outstanding memorization capacity and exposure to KITTI during training. Because overall noise resilience on RobustSpring qualitatively transfers to KITTI, RobustSpring supports model selection for real-world settings where corruption robustness cannot be measured.

## 5 CONCLUSION

With RobustSpring we introduce an image corruption dataset and benchmark that evaluates the robustness of optical flow, scene flow, and stereo models. We carefully design 20 different image corruptions and integrate them in time, stereo, and depth for a holistic evaluation of dense matching tasks. Furthermore, we establish a corruption robustness metric using clean and corrupted predictions, and compare ranking strategies to unify model results across all 20 corruptions. RobustSpring’s benchmark further supports data-efficient result uploads to an evaluation server. Our initial evaluation of 17 optical flow, scene flow, and stereo models reveals an overall high sensitivity to corrupted images. As our robustness results translate to real-world performance, systematic corruption benchmarks like RobustSpring are crucial to uncover potential model performance improvements.

**Limitations.** Due to its benchmark character, we have limited the image corruptions on RobustSpring to a selection of 20. While this does not cover the full space of potential corruptions, this data-budget limitation is necessary to make the RobustSpring dataset applicable and not overburden the computational resources of researchers during evaluation.

540 REFERENCES  
541

542 Shashank Agnihotri, Kanchana Vaishnavi Gandikota, Julia Grabinski, Paramanand Chandramouli,  
543 and Margret Keuper. On the unreasonable vulnerability of transformers for image restoration-  
544 and an easy fix. In *Proc. IEEE/CVF International Conference on Computer Vision Workshops*  
545 (*ICCVW*), 2023.

546 Shashank Agnihotri, Julia Grabinski, Janis Keuper, and Margret Keuper. Beware of aliases—signal  
547 preservation is crucial for robust image restoration. *arXiv preprint 2406.07435*, 2024a.

548 Shashank Agnihotri, Julia Grabinski, and Margret Keuper. Improving feature stability during up-  
549 sampling – spectral artifacts and the importance of spatial context. In *Proc. European Conference*  
550 *on Computer Vision (ECCV)*, pp. 357–376, 2024b.

551 Shashank Agnihotri, Steffen Jung, and Margret Keuper. CosPGD: an efficient white-box adversarial  
552 attack for pixel-wise prediction tasks. In *Proc. International Conference on Learning Represen-*  
553 *tations (ICML)*, pp. 416–451, 2024c.

554 Adithya Prem Anand, H. Gokul, Harish Srinivasan, Pranav Vijay, and Vineeth Vijayaraghavan. Ad-  
555 versarial patch defense for optical flow networks in video action recognition. In *Proc. IEEE In-*  
556 *ternational Conference on Machine Learning and Applications (ICMLA)*, pp. 1289–1296, 2020.

557 Simon Baker, Daniel Scharstein, J. P. Lewis, Stefan Roth, Michael J. Black, and Richard Szeliski. A  
558 database and evaluation methodology for optical flow. *International Journal of Computer Vision*  
559 (*IJCV*), 92(1):1–31, 2011.

560 Zachary Berger, Parth Agrawal, Tyan Yu Liu, Stefano Soatto, and Alex Wong. Stereoscopic univer-  
561 *562* sal perturbations across different architectures and datasets. In *Proc. IEEE/CVF Conference on*  
563 *Computer Vision and Pattern Recognition (CVPR)*, pp. 15180–15190, 2022.

564 Daniel J. Butler, Jonas Wulff, G. B. Stanley, and Michael J. Black. A naturalistic open source  
565 movie for optical flow evaluation. In *Proc. European Conference on Computer Vision (ECCV)*,  
566 pp. 611–625, 2012.

567 Xuelian Cheng, Yiran Zhong, Mehrtash Harandi, Yuchao Dai, Xiaojun Chang, Hongdong Li, Tom  
568 Drummond, and Zongyuan Ge. Hierarchical neural architecture search for deep stereo matching.  
569 *Proc. Conference on Neural Information Processing Systems (NeurIPS)*, 33:22158–22169, 2020.

570 Francesco Croce, Maksym Andriushchenko, Vikash Sehwag, Edoardo Debenedetti, Nicolas Flam-  
571 marion, Mung Chiang, Prateek Mittal, and Matthias Hein. RobustBench: a standardized adver-  
572 *573* sarial robustness benchmark. In *Proc. Conference on Neural Information Processing Systems*  
574 (*NeurIPS*), 2021.

575 Alexey Dosovitskiy, Philipp Fischer, Eddy Ilg, Caner Hazirbas, Vladimir Golkov, Patrick van der  
576 Smagt, Daniel Cremers, and Thomas Brox. FlowNet: learning optical flow with convolutional  
577 networks. In *Proc. IEEE/CVF International Conference on Computer Vision (ICCV)*, 2015.

578 Nathan Drenkow and Mathias Unberath. RobustCLEVR: A benchmark and framework for evaluat-  
579 ing robustness in object-centric learning. In *Proc. IEEE/CVF Winter Conference on Applications*  
580 *581* *of Computer Vision (WACV)*, pp. 4518–4527, 2024.

582 Adrien Gaidon, Qiao Wang, Yohann Cabon, and Eleonora Vig. Virtual worlds as proxy for multi-  
583 *584* object tracking analysis. In *Proc. IEEE/CVF Conference on Computer Vision and Pattern Recog-*  
585 *586* *nition (CVPR)*, pp. 4340–4349, 2016.

587 Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for autonomous driving? the KITTI  
588 vision benchmark suite. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition*  
589 (*CVPR*), pp. 3354–3361, 2012.

590 Shirsendu Sukanta Halder, Jean-François Lalonde, and Raoul de Charette. Physics-based rendering  
591 for improving robustness to rain. In *Proc. IEEE/CVF Conference on Computer Vision and Pattern*  
592 *593* *Recognition (CVPR)*, pp. 10203–10212, 2019.

594 Matthias Hein and Maksym Andriushchenko. Formal guarantees on the robustness of a classifier  
 595 against adversarial manipulation. In *Proc. Conference on Neural Information Processing Systems*  
 596 (*NIPS*), volume 30, 2017.

597

598 Dan Hendrycks and Thomas Dietterich. Benchmarking neural network robustness to common  
 599 corruptions and perturbations. In *Proc. International Conference on Learning Representations*  
 600 (*ICLR*), pp. 1–16, 2019.

601

602 Alain Hore and Djemel Ziou. Image quality metrics: PSNR vs. SSIM. In *Proc. IEEE International*  
 603 *Conference on Pattern Recognition (ICPR)*, pp. 2366–2369, 2010.

604

605 Berthold KP Horn and Brian G Schunck. Determining optical flow. *Artificial intelligence*, 17(1-3):  
 185–203, 1981.

606

607 Zhaoyang Huang, Xiaoyu Shi, Chao Zhang, Qiang Wang, Ka Chun Cheung, Hongwei Qin, Jifeng  
 608 Dai, and Hongsheng Li. FlowFormer: a transformer architecture for optical flow. In *Proc. Euro-*  
 609 *pean Conference on Computer Vision (ECCV)*, pp. 668–685, 2022.

610

611 Eddy Ilg, Nikolaus Mayer, Tonmoy Saikia, Margret Keuper, Alexey Dosovitskiy, and Thomas Brox.  
 612 FlowNet 2.0: Evolution of optical flow estimation with deep networks. In *Proc. IEEE/CVF Con-*  
 613 *ference on Computer Vision and Pattern Recognition (CVPR)*, pp. 2462–2470, 2017.

614

615 John Immerkaer. Fast noise variance estimation. *Computer Vision and Image Understanding*  
 (CVIU), 64(2):300–302, 1996.

616

617 Azin Jahedi, Lukas Mehl, Marc Rivinius, and Andrés Bruhn. Multi-scale RAFT: Combining hier-  
 618 archical concepts for learning-based optical flow estimation. In *Proc. IEEE International Confer-*  
 619 *ence on Image Processing (ICIP)*, pp. 1236–1240, 2022.

620

621 Azin Jahedi, Maximilian Luz, Marc Rivinius, Lukas Mehl, and Andrés Bruhn. MS-RAFT+: High  
 622 resolution multi-scale RAFT. *International Journal of Computer Vision (IJCV)*, 132(5):1835–  
 623 1856, 2024.

624

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

628

629 Steffen Jung, Jovita Lukasik, and Margret Keuper. Neural architecture design and robustness: A  
 630 dataset. In *Proc. International Conference on Learning Representations (ICLR)*. OpenReview.  
 net, 2023.

631

632 Oğuzhan Fatih Kar, Teresa Yeo, Andrei Atanov, and Amir Zamir. 3D common corruptions and  
 633 data augmentation. In *Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition*  
 (CVPR), pp. 18963–18974, 2022.

634

635 Daniel Kondermann, Rahul Nair, Katrin Honauer, Karsten Krispin, Jonas Andrulis, Alexander  
 636 Brock, Burkhard Gussfeld, Mohsen Rahimimoghaddam, Sabine Hofmann, Claus Brenner, et al.  
 637 The HCI benchmark suite: Stereo and flow ground truth with uncertainties for urban autonomous  
 638 driving. In *Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*  
 (CVPRW), pp. 19–28, 2016.

639

640 Lingdong Kong, Youquan Liu, Xin Li, Runnan Chen, Wenwei Zhang, Jiawei Ren, Liang Pan, Kai  
 641 Chen, and Ziwei Liu. Robo3D: Towards robust and reliable 3d perception against corruptions. In  
 642 *Proc. IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 19994–20006, 2023.

643

644 Tom Koren, Lior Talker, Michael Dinerstein, and Ran Vitek. Consistent semantic attacks on optical  
 645 flow. In *Proc. Asian Conference on Computer Vision (ACCV)*, pp. 1658–1674, 2022.

646

647 Lorenzo Lamberti, Lorenzo Bellone, Luka Macan, Enrico Natalizio, Francesco Conti, Daniele  
 648 Palossi, and Luca Benini. Distilling tiny and ultra-fast deep neural networks for autonomous  
 649 navigation on nano-uavs. *IEEE Internet of Things Journal (IOTJ)*, 2024.

648 Yijin Li, Yichen Shen, Zhaoyang Huang, Shuo Chen, Weikang Bian, Xiaoyu Shi, Fu-Yun Wang,  
 649 Keqiang Sun, Hujun Bao, Zhaopeng Cui, Guofeng Zhang, and Hongsheng Li. BlinkVision: A  
 650 benchmark for optical flow, scene flow and point tracking estimation using rgb frames and events.  
 651 *arXiv preprint 2410.20451*, 2024.

652 Lahav Lipson, Zachary Teed, and Jia Deng. RAFT-stereo: Multilevel recurrent field transforms for  
 653 stereo matching. In *Proc. International Conference on 3D Vision (3DV)*, pp. 218–227, 2021.

654 K. T. Yasar Mahima, Asanka G. Perera, Sreenatha Anavatti, and Matt Garratt. FlowCraft: Unveiling  
 655 adversarial robustness of LiDAR scene flow estimation. *Pattern Recognition Letters (PRL)*, 2025.

656 Daniel Maurer, Nico Marniok, Bastian Goldluecke, and Andrés Bruhn. Structure-from-motion-  
 657 aware patchmatch for adaptive optical flow estimation. In *Proc. European Conference on Com-  
 658 puter Vision (ECCV)*, pp. 565–581, 2018.

659 Nikolaus Mayer, Eddy Ilg, Philip Hausser, Philipp Fischer, Daniel Cremers, Alexey Dosovitskiy,  
 660 and Thomas Brox. A large dataset to train convolutional networks for disparity, optical flow,  
 661 and scene flow estimation. In *Proc. IEEE/CVF Conference on Computer Vision and Pattern  
 662 Recognition (CVPR)*, pp. 4040–4048, 2016.

663 Kimberly McGuire, Guido De Croon, Christophe De Wagter, Karl Tuyls, and Hilbert Kappen. Ef-  
 664 ficient optical flow and stereo vision for velocity estimation and obstacle avoidance on an au-  
 665 tonomous pocket drone. *IEEE Robotics and Automation Letters (RA-L)*, 2(2):1070–1076, 2017.

666 Morgan McGuire and Louis Bavoil. Weighted blended order-independent transparency. *Journal of  
 667 Computer Graphics Techniques (JCGT)*, 2013.

668 Lukas Mehl, Azin Jahedi, Jenny Schmalfuss, and Andrés Bruhn. M-FUSE: multi-frame fusion for  
 669 scene flow estimation. In *Proc. IEEE/CVF Winter Conference on Applications of Computer Vision  
 670 (WACV)*, pp. 2020–2029, 2023a.

671 Lukas Mehl, Jenny Schmalfuss, Azin Jahedi, Yaroslava Nalivayko, and Andrés Bruhn. Spring: A  
 672 high-resolution high-detail dataset and benchmark for scene flow, optical flow and stereo. In  
 673 *Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 4981–  
 674 4991, 2023b.

675 Lukas Mehl, Andrés Bruhn, Markus Gross, and Christopher Schroers. Stereo conversion with  
 676 disparity-aware warping, compositing and inpainting. In *Proc. IEEE/CVF Winter Conference  
 677 on Applications of Computer Vision (WACV)*, pp. 4260–4269, 2024.

678 Moritz Menze and Andreas Geiger. Object scene flow for autonomous vehicles. In *Proc. IEEE/CVF  
 679 Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 3061–3070, 2015.

680 Claudio Michaelis, Benjamin Mitzkus, Robert Geirhos, Evgenia Rusak, Oliver Bringmann, Alexander  
 681 S. Ecker, Matthias Bethge, and Wieland Brendel. Benchmarking robustness in object detec-  
 682 tion: Autonomous driving when winter is coming. In *Proc. Conference on Neural Information  
 683 Processing Systems Workshops (NeurIPS-Workshops)*, 2019.

684 Sergiu Mocanu, Alan R Moody, and April Khademi. FlowReg: fast deformable unsupervised med-  
 685 ical image registration using optical flow. *Journal of Machine Learning for Biomedical Imaging  
 686 (MELBA)*, 2021.

687 Patrick Müller, Alexander Braun, and Margret Keuper. Classification robustness to common op-  
 688 tical aberrations. In *Proc. IEEE/CVF International Conference on Computer Vision Workshops  
 689 (ICCVW)*, pp. 3632–3643, 2023.

690 Patricia Pauli, Anne Koch, Julian Berberich, Paul Kohler, and Frank Allgöwer. Training robust  
 691 neural networks using Lipschitz bounds. In *IEEE Control Systems Letters (L-CSS)*, volume 6, pp.  
 692 121–126, 2022.

693 Tan-Binh Phan, Dinh-Hoan Trinh, Didier Wolf, and Christian Daul. Optical flow-based structure-  
 694 from-motion for the reconstruction of epithelial surfaces. *Pattern Recognition (PR)*, 2020.

702 Markus Philipp, Neal Bacher, Stefan Saur, Franziska Mathis-Ullrich, and Andrés Bruhn. From  
 703 chairs to brains: customizing optical flow for surgical activity localization. In *Proc. IEEE International*  
 704 *Symposium on Biomedical Imaging (ISBI)*, 2022.

705 Anurag Ranjan and Michael J Black. Optical flow estimation using a spatial pyramid network. In  
 706 *Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 4161–  
 707 4170, 2017.

708 Anurag Ranjan, Joel Janai, Andreas Geiger, and Michael J. Black. Attacking optical flow. In *Proc.*  
 709 *IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 2004–2013, 2019.

710 Anurag Ranjan, David T. Hoffmann, Dimitrios Tzionas, Siyu Tang, Javier Romero, and Michael J.  
 711 Black. Learning multi-human optical flow. *International Journal of Computer Vision (IJCV)*, 128  
 712 (4):873–890, 2020.

713 Stephan Richter, Zeeshan Hayder, and Vladlen Koltun. Playing for benchmarks. In *Proc. IEEE/CVF*  
 714 *International Conference on Computer Vision (ICCV)*, pp. 2232–2241, 2017.

715 Benoît Rosa, Valentin Bordoux, and Florent Nageotte. Combining differential kinematics and op-  
 716 tical flow for automatic labeling of continuum robots in minimally invasive surgery. *Frontiers in*  
 717 *Robotics and AI*, 6:86, 2019.

718 Daniel Scharstein, Heiko Hirschmüller, York Kitajima, Greg Krathwohl, Nera Nesic, Xi Wang, and  
 719 Porter Westling. High-resolution stereo datasets with subpixel-accurate ground truth. In *Proc.*  
 720 *German Conference on Pattern Recognition (GCPR)*, 2014.

721 Erik Scheurer, Jenny Schmalfuss, Alexander Lis, and Andrés Bruhn. Detection defenses: An empty  
 722 promise against adversarial patch attacks on optical flow. In *Proc. IEEE/CVF Winter Conference*  
 723 *on Applications of Computer Vision (WACV)*, pp. 6489–6498, 2024.

724 Jenny Schmalfuss, Lukas Mehl, and Andrés Bruhn. Attacking motion estimation with adversarial  
 725 snow. *ECCV 2022 Workshop on Adversarial Robustness in the Real World (ECCV-AROW)*, 2022a.

726 Jenny Schmalfuss, Philipp Scholze, and Andrés Bruhn. A perturbation-constrained adversarial at-  
 727 tack for evaluating the robustness of optical flow. In *Proc. European Conference on Computer*  
 728 *Vision (ECCV)*, pp. 183–200, 2022b.

729 Jenny Schmalfuss, Lukas Mehl, and Andrés Bruhn. Distracting downpour: Adversarial weather  
 730 attacks for motion estimation. In *Proc. IEEE/CVF International Conference on Computer Vision*  
 731 *(ICCV)*, pp. 10106–10116, 2023.

732 Thomas Schöps, Johannes L. Schönberger, Silvano Galliani, Torsten Sattler, Konrad Schindler, Marc  
 733 Pollefeys, and Andreas Geiger. A multi-view stereo benchmark with high-resolution images and  
 734 multi-camera videos. In *Proc. IEEE/CVF Conference on Computer Vision and Pattern Recog-*  
 735 *nition (CVPR)*, 2017.

736 Simon Schrodi, Tonmoy Saikia, and Thomas Brox. Towards understanding adversarial robustness  
 737 of optical flow networks. In *Proc. IEEE/CVF Conference on Computer Vision and Pattern Recog-*  
 738 *nition (CVPR)*, pp. 8916–8924, 2022.

739 Markus Schulze. The Schulze method of voting. *arXiv preprint arXiv:1804.02973*, 2018.

740 Deqing Sun, Xiaodong Yang, Ming-Yu Liu, and Jan Kautz. PWC-Net: CNNs for optical flow using  
 741 pyramid, warping, and cost volume. In *Proc. IEEE/CVF Conference on Computer Vision and*  
 742 *Pattern Recognition (CVPR)*, 2018.

743 Deqing Sun, Daniel Vlasic, Charles Herrmann, Varun Jampani, Michael Krainin, Huiwen Chang,  
 744 Ramin Zabih, William T. Freeman, , and Ce Liu. AutoFlow: Learning a better training set for op-  
 745 tical flow. In *Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*,  
 746 2021.

747 Shiyu Tang, Ruihao Gong, Yan Wang, Aishan Liu, Jiakai Wang, Xinyun Chen, Fengwei Yu, Xi-  
 748 anglong Liu, Dawn Song, Alan Yuille, Philip H.S. Torr, and Dacheng Tao. RobustART: Bench-  
 749 marking robustness on architecture design and training techniques. *arXiv preprint 2109.05211*,  
 750 2021.

751

756 Rohan Taori, Achal Dave, Vaishaal Shankar, Nicholas Carlini, Benjamin Recht, and Ludwig  
 757 Schmidt. Measuring robustness to natural distribution shifts in image classification. In *Proc.*  
 758 *Conference on Neural Information Processing Systems (NeurIPS)*, volume 33, pp. 18583–18599,  
 759 2020.

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

762 Zachary Teed and Jia Deng. RAFT-3D: Scene flow using rigid-motion embeddings. In *Proc.*  
 763 *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 8375–8384,  
 764 2021.

765 Dimitris Tsipras, Shibani Santurkar, Logan Engstrom, Alexander Turner, and Aleksander Madry.  
 766 Robustness may be at odds with accuracy. In *International Conference on Learning Representations (ICLR)*, 2019. URL <https://openreview.net/forum?id=SyxAb30cY7>.

767 Pengfei Wang, Xiaofei Hui, Beijia Lu, Nimrod Lilith, Jun Liu, and Sameer Alam. Left-right dis-  
 768 crepancy for adversarial attack on stereo networks. *arXiv preprint 2401.07188*, 2024a.

769 Yihan Wang, Lahav Lipson, and Jia Deng. Sea-raft: Simple, efficient, accurate raft for optical flow.  
 770 In *European Conference on Computer Vision*, pp. 36–54. Springer, 2024b.

771 Thomas Wiesemann and Xiaoyi Jiang. Fog augmentation of road images for performance analysis  
 772 of traffic sign detection algorithms. In *Proc. International Conference on Advanced Concepts for*  
 773 *Intelligent Vision Systems (ACVIS)*, 2016.

774 Alex Wong, Mukund Mundhra, and Stefano Soatto. Stereopagnosia: Fooling stereo networks with  
 775 adversarial perturbations. *Proc. AAAI Conference on Artificial Intelligence (AAAI)*, pp. 2879–  
 776 2888, 2021.

777 Gangwei Xu, Junda Cheng, Peng Guo, and Xin Yang. Attention concatenation volume for accurate  
 778 and efficient stereo matching. In *Proc. IEEE/CVF Conference on Computer Vision and Pattern*  
 779 *Recognition (CVPR)*, pp. 12981–12990, 2022a.

780 Haofei Xu, Jing Zhang, Jianfei Cai, Hamid Rezatofighi, and Dacheng Tao. Gmflow: Learning  
 781 optical flow via global matching. In *Proc. IEEE/CVF Conference on Computer Vision and Pattern*  
 782 *Recognition (CVPR)*, pp. 8121–8130, 2022b.

783 Koichiro Yamanaka, Keita Takahashi, Toshiaki Fujii, and Ryuraroh Matsumoto. Simultaneous at-  
 784 tack on CNN-based monocular depth estimation and optical flow estimation. *IEICE Transactions*  
 785 *on Information and Systems*, pp. 785–788, 2021.

786 Zhonghua Yi, Hao Shi, Qi Jiang, Yao Gao, Ze Wang, Yufan Zhang, Kailun Yang, and Kaiwei  
 787 Wang. Benchmarking the robustness of optical flow estimation to corruptions. *arXiv preprint*  
 788 *2411.14865*, 2024.

789 Feihu Zhang, Victor Prisacariu, Ruigang Yang, and Philip HS Torr. GA-Net: guided aggregation net  
 790 for end-to-end stereo matching. In *Proc. IEEE/CVF Conference on Computer Vision and Pattern*  
 791 *Recognition (CVPR)*, pp. 185–194, 2019.

792 Qingwen Zhang, Yi Yang, Peizheng Li, Olov Andersson, and Patric Jensfelt. SeFlow: a self-  
 793 supervised scene flow method in autonomous driving. In *Proc. European Conference on Com-  
 794 puter Vision (ECCV)*, pp. 353–369, 2025.

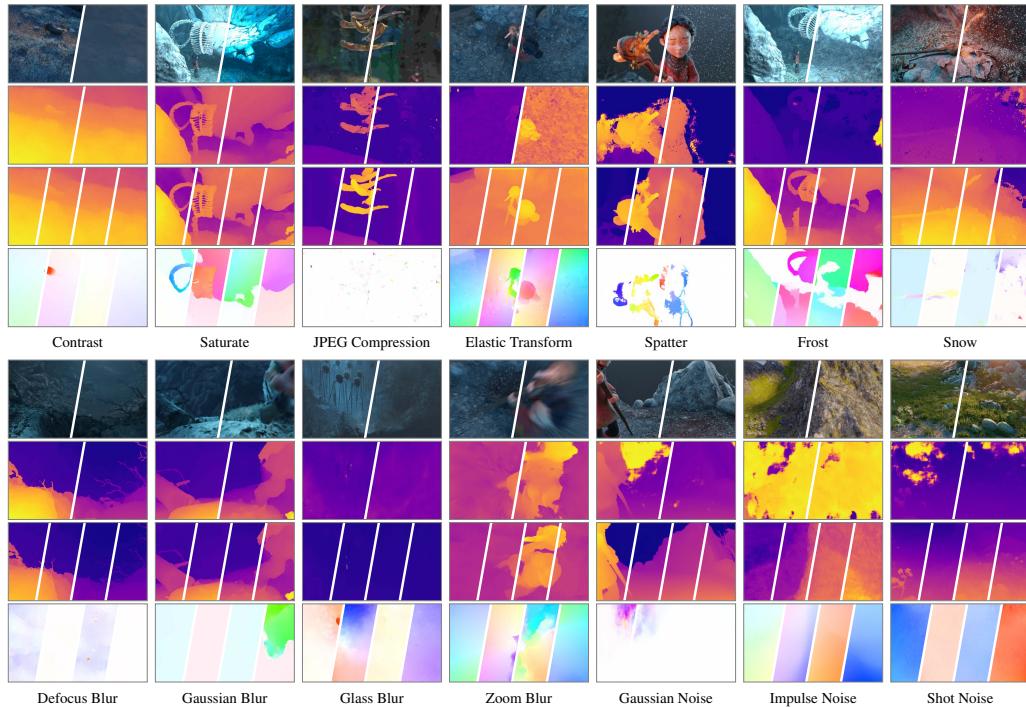
795 Yinjiang Zheng, Mingfang Zhang, and Feng Lu. Optical flow in the dark. In *Proc. IEEE/CVF*  
 796 *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.

797 Yuanhang Zheng, Harald Köstler, Nils Thürey, and Ulrich Rüde. Enhanced motion blur calculation  
 798 with optical flow. In *Proc. Workshop on Vision, Modeling and Visualization (VMV)*, 2006.

810 A APPENDIX  
811812 A.1 DISCUSSION OF ROBUSTNESS DEFINITION  
813814 **Scope of RobustSpring’s robustness definition.** Robustness is a multifaceted concept. In Robust-  
815 Spring we focus on one specific and mathematically established perspective: robustness as predic-  
816 tion stability under input corruptions. This choice is motivated by its independence from ground-  
817 truth definitions of optical flow, and by its practical utility in assessing whether models maintain  
818 consistent scene estimates in the presence of visual disturbances. Other notions of robustness, such  
819 as those directly tied to ground truth, are equally valid but outside the scope of this benchmark.820 **Accuracy vs. Robustness.** RobustSpring explicitly disentangles *accuracy* and *robustness*. Accuracy  
821 measures the deviation between model predictions and ground truth, which depends on the adopted  
822 definition of optical flow—here, the true 3D motion of visible surfaces projected into the 2D image  
823 plane (Horn & Schunck, 1981; Baker et al., 2011). Robustness, by contrast, is defined via Lipschitz  
824 continuity (Eq. 1) as the stability of predictions under corrupted inputs, independent of ground truth.825 **Why Ground-truth-free Robustness.** Some works define robustness as the difference between  
826 predictions on corrupted images and the corresponding ground truth (Ranjan et al., 2019; Agnihotri  
827 et al., 2024c). This approach entangles robustness with accuracy and requires ambiguous judgments  
828 about how each corruption changes the motion field (*e.g.* elastic transforms, rain, snow). By adopt-  
829 ing a ground-truth-free definition, RobustSpring avoids this ambiguity and aligns with established  
830 robustness literature (Schmalfuss et al., 2022b; Tsipras et al., 2019; Taori et al., 2020; Hein & An-  
831 driushchenko, 2017; Pauli et al., 2022).832 **Perfect Accuracy vs. Perfect Robustness.** Under our definitions, a perfectly accurate method may  
833 achieve poor robustness scores for corruptions that genuinely alter the motion field (*e.g.* snow, elas-  
834 tic transform), because its predictions change in accordance with the altered scene. Conversely, a  
835 perfectly robust method (constant predictions) is maximally stable but inaccurate. This disentan-  
836 glement is intentional: it reveals cases where accuracy and robustness align versus diverge. Such  
837 divergence highlights model behavior that plain accuracy metrics cannot capture.838 **Practical Relevance.** As demonstrated in Fig. 5a, current models tend to propagate spurious fore-  
839 ground motions introduced by rain into the background, leading to degraded scene estimates. Ro-  
840 bustSpring’s robustness metric captures this instability and thus provides valuable insight into real-  
841 world performance. The combination of accuracy and robustness offers practitioners a two-axis  
842 evaluation, allowing trade-offs between precision and stability depending on the target application.844 A.2 LINKS AND CHECKPOINTS FOR EVALUATED MODELS  
845846 We evaluated the original, author provided optical flow, scene flow and stereo methods on the Ro-  
847 bustSpring dataset. Tab. 4 reports the repositories and checkpoints for the optical flow, scene flow  
848 and stereo models, which were benchmarked on RobustSpring in Tab. 1, Tab. 2a, and Tab. 2b. Fur-  
849 ther details on training and checkpoints for these models can be found in their original publications.850 **Ressources and Run Times.** We conducted all evaluations on a single NVIDIA RTX A6000 (48  
851 GB) GPU. The evaluation time of models on the RobustSpring data strongly depends on the compu-  
852 tational efficiency and requirements of the original models. As a representative example, the optical  
853 flow evaluation with RAFT (Teed & Deng, 2020) took 20 hours on the full RobustSpring data.854 Once the methods are evaluated and the results uploaded to our benchmark server prototype, the ro-  
855 bustness evaluation for an optical flow method takes 13 minutes on our evaluation machine (Intel(R)  
856 Xeon(R) Gold 6130 CPU @ 2.10GHz, 8 vCPUs). The robustness evaluation for scene flow takes  
857 about 26 minutes, and 7 minutes for stereo.859 A.3 ROBUSTSPRING IMAGE CORRUPTIONS  
860861 Below, we provide supplementary information on the image corruptions for the RobustSpring  
862 dataset. Besides visualizing further benchmark samples and supplying a video that showcases the  
863 space- and time-integration of our corruptions, we also give details on their implementation with a  
focus on RobustSpring-specific consistencies.

864  
 865  
 866  
 867  
 868  
 869  
 Table 4: Repositories and checkpoints used for evaluating methods in RobustSpring. \*The Mixed  
 870 checkpoint MS-RAFT+ is pretrained on Chairs and Things and then fine-tuned on a mix of Sintel,  
 871 Viper, KITTI 2015, Things, and HD1k.  
 872  
 873  
 874  
 875  
 876

Method	Repository	Checkpoint
<b>Optical Flow</b>		
RAFT	<a href="https://github.com/princeton-vl/RAFT">https://github.com/princeton-vl/RAFT</a>	Sintel
PWCNet	<a href="https://github.com/NVlabs/PWC-Net">https://github.com/NVlabs/PWC-Net</a>	Sintel
GMFlow	<a href="https://github.com/haofeixu/gmflow">https://github.com/haofeixu/gmflow</a>	Sintel
GMA	<a href="https://github.com/zacjiang/GMA">https://github.com/zacjiang/GMA</a>	Sintel
FlowNet2	<a href="https://github.com/NVIDIA/flownet2-pytorch">https://github.com/NVIDIA/flownet2-pytorch</a>	Sintel
FlowFormer	<a href="https://github.com/drinkingcoder/FlowFormer-Official">https://github.com/drinkingcoder/FlowFormer-Official</a>	Sintel
MS-RAFT+	<a href="https://github.com/cv-stuttgart/MS_RAFT_plus">https://github.com/cv-stuttgart/MS_RAFT_plus</a>	Mixed*
SEA-RAFT	<a href="https://github.com/princeton-vl/SEA-RAFT">https://github.com/princeton-vl/SEA-RAFT</a>	Spring (M)
SPyNet	<a href="https://github.com/anuragranj/floattack (PyTorch implementation)">https://github.com/anuragranj/floattack (PyTorch implementation)</a>	Sintel
<a href="https://github.com/anuragranj/spynet (original implementation)">https://github.com/anuragranj/spynet (original implementation)</a>		
<b>Scene Flow</b>		
M-FUSE	<a href="https://github.com/cv-stuttgart/M-FUSE">https://github.com/cv-stuttgart/M-FUSE</a>	KITTI 2015
<b>Stereo</b>		
RAFT-Stereo	<a href="https://github.com/princeton-vl/RAFT-Stereo/">https://github.com/princeton-vl/RAFT-Stereo/</a>	Scene Flow (s)
ACVNet	<a href="https://github.com/gangweiX/ACVNet">https://github.com/gangweiX/ACVNet</a>	Scene Flow (s)
LEAStereo	<a href="https://github.com/XuelianCheng/LEAStereo">https://github.com/XuelianCheng/LEAStereo</a>	Scene Flow (s), KITTI (k)
GANet	<a href="https://github.com/feihuzhang/GANet">https://github.com/feihuzhang/GANet</a>	Scene Flow (s), KITTI (k)



907  
 908  
 909  
 910  
 911  
 912  
 913  
 914  
 915  
 916  
 917  
 Figure 6: RobustSpring example frames, complementing Fig. 3. The first row shows clean and  
 918 corrupted images. The second row shows the left and right disparity maps predicted with LEA  
 919 Stereo (Cheng et al., 2020). The third row shows the target disparities for forward left, backward  
 920 left, forward right, and backward right directions from M-FUSE (Mehl et al., 2023a). The fourth  
 921 row shows optical flow estimates for forward left, backward left, forward right, and backward right  
 922 from RAFT (Teed & Deng, 2020). All disparities and flows are computed on the corrupted dataset.

### A.3.1 ADDITIONAL CORRUPTION BENCHMARK SAMPLES

923  
 924  
 925  
 926  
 927  
 To complement the benchmark samples in Fig. 3, we show benchmark results on additional corrup-  
 928 tions in Fig. 6.

918 A.3.2 VIDEO OF CORRUPTIONS  
919920 With the supplementary material, we include a video that visualizes all corruptions applied to the  
921 Spring test dataset.  
922923 A.3.3 CORRUPTION IMPLEMENTATION  
924925 Below we provide the implementation details and parameters for all corruptions in the RobustSpring  
926 dataset. We cluster the corruptions by their main classes. The original (uncorrupted) image is  
927 denoted as  $I$ , while the corrupted version is  $\hat{I}$ . The pixel-aligned depth values are  $D$ . In the stereo  
928 video setting, the image subscripts  $t$  and  $t + 1$  denote frames over time, while  $l$  and  $r$  denote left  
929 and right frame, where necessary. All of RobustSpring’s noises are independent of time, stereo and  
930 depth, which means they are sampled independently for every single image of the dataset.  
931**Brightness.** The brightness is adapted via

932 
$$\hat{I} = I + c, \quad (4)$$
  
933

934 and for time- and stereo-consistent brightness changes in RobustSpring we choose the parameter  
935  $c = c_t^l = c_t^r = c_{t+1}^l = c_{t+1}^r = 0.39$ .  
936**Contrast.** The equation to adapt contrast is

937 
$$\hat{I} = (I - \text{mean}(I)) \cdot c + \text{mean}(I), \quad (5)$$
  
938

939 where we selected  $c = c_t^l = c_t^r = c_{t+1}^l = c_{t+1}^r = 0.16$  for time- and stereo-consistent contrast  
940 adaptations.  
941**Saturation.** For those adaptations the RGB image is transformed to HSV, and the saturation component  $S$  is adapted via

942 
$$\hat{S} = S \cdot \alpha + \beta, \quad (6)$$
  
943

944 with  $\alpha = \alpha_t^l = \alpha_t^r = \alpha_{t+1}^l = \alpha_{t+1}^r = 2.3$  and  $\beta_t^l = \beta_t^r = \beta_{t+1}^l = \beta_{t+1}^r = 0.01$  for time- and  
945 stereo-consistent saturation changes.  
946**Defocus Blur.** The defocus blur convolves the image with a circular mean filter  $C^{mean}$ 

947 
$$\hat{I} = I * C_r^{mean}, \quad (7)$$
  
948

949 where we choose the radius  $r = r_t^l = r_t^r = r_{t+1}^l = r_{t+1}^r = 6$  for time- and stereo-consistent  
950 blurring.  
951**Gaussian Blur.** Gaussian blur convolves the image with a Gaussian  $C^{gauss}$ 

952 
$$\hat{I} = I * C_\sigma^{gauss}, \quad (8)$$
  
953

954 where we choose the standard deviation  $\sigma = \sigma_t^l = \sigma_t^r = \sigma_{t+1}^l = \sigma_{t+1}^r = 4$  for time- and stereo-  
955 consistent blurring.  
956**Glass Blur.** This is a Gauss-blurred image, whose pixels are afterwards shuffled via the shuffling  
957  $S(I, i, r)$  over several iterations  $i$  within a neighborhood of radius  $r$   
958

959 
$$\hat{I} = S(I * C_\sigma^{gauss}, i, r), \quad (9)$$
  
960

961 where two sets of time-consistent parameters are picked for the different stereo cameras:  $\sigma_t^l =$   
962  $\sigma_{t+1}^l = 1.2$ ,  $\sigma_t^r = \sigma_{t+1}^r = 1.2$ ,  $i_t^l = i_{t+1}^l = 1$ ,  $i_t^r = i_{t+1}^r = 1$ ,  $r_t^l = r_{t+1}^l = 3$  and  $r_t^r = r_{t+1}^r = 3$ .  
963**Motion Blur.** Motion blur is implemented by averaging the intensities of pixels along the motion  
964 trajectory determined by the optical flow. Let  $\mathbf{v}(x, y) = (v_x(x, y), v_y(x, y))$  be the optical flow  
965 vector at pixel  $(x, y)$ , and let

966 
$$N = \max \left( 1, \left\lfloor 10 \cdot \max_{(x,y)} \|\mathbf{v}(x, y)\| \right\rfloor \right) \quad (10)$$
  
967

968 be the number of samples along the motion path. Then, the blurred pixel is computed as  
969

970 
$$\hat{I}(x, y) = \frac{1}{N+1} \sum_{k=0}^N I \left( x + \frac{k}{N} v_x(x, y), y + \frac{k}{N} v_y(x, y) \right). \quad (11)$$
  
971

972 Here, the scaling factor 10 controls the extent of the blur relative to the magnitude of the motion.  
 973

974 **Zoom Blur.** Zoom blur is created by averaging the original image with a series of zoomed-in  
 975 versions of itself. Specifically, let  $Z(I, z)$  denote the image  $I$  zoomed by a factor  $z$ , and let  $\{z_i\}$  be  
 976 a set of zoom factors ranging from 1 to approximately 1.24 (in increments of 0.02). Then the final  
 977 image is computed as

$$978 \hat{I} = \frac{1}{N+1} \left( I + \sum_{i=1}^N Z(I, z_i) \right), \quad (12)$$

980 where  $N$  is the number of zoom factors. This formulation averages the original image with its  
 981 progressively zoomed versions, resulting in a smooth zoom blur effect.  
 982

983 **Gaussian Noise.** This noise adds a random value from a Normal distribution to every pixel in the  
 984 original image, where  $\mathcal{N}_I(\mu, \sigma^2)$  is a  $I$ -shaped array of random numbers that are drawn from the  
 985 Normal distribution with mean  $\mu$  and variance  $\sigma^2$ :

$$986 \hat{I} = I + \alpha \cdot \mathcal{N}_I(0, 1). \quad (13)$$

987 The scaling  $\alpha = 0.115$  is selected for all images in RobustSpring, but  $\mathcal{N}_I(0, 1)$  is sampled anew for  
 988 every image.  
 989

990 **Impulse Noise.** Here, for a fixed fraction of pixels  $p$ , their values are replaced by the values 0 or  
 991 255. For RobustSpring,  $p = 0.075$ .

992 **Speckle Noise.** Like Gaussian noise, Speckle noise also builds on random values from a Normal  
 993 distribution, but adds these values after additionally scaling with  $I$ :

$$994 \hat{I} = I + I \cdot \alpha \cdot \mathcal{N}_I(0, 1). \quad (14)$$

995 For RobustSpring, the parameter is  $\alpha = 0.45$ .

996 **Shot Noise.** Shot noise uses values drawn from a Poisson distribution  $\mathcal{P}$  per pixel

$$997 \hat{I} = \frac{\mathcal{P}(I \cdot c)}{c}, \quad (15)$$

1001 where  $c = 23$  for RobustSpring.  
 1002

1003 **Pixelation.** This is achieved by downsampling the image to size  $s$ , a fraction of its original size,  
 1004 with a box filter  $box(I, s)$ , followed by upsampling  $up(I, s)$  to the size  $s$ , which is the original size:  
 1005

$$1006 \hat{I} = up(box(I, I \cdot c), I). \quad (16)$$

1007 For RobustSpring, we use the size fraction  $c = c_t^l = c_t^r = c_{t+1}^l = c_{t+1}^r = 0.16$  for time- and  
 1008 stereo-consistent pixelation.  
 1009

1010 **JPEG Compression.** For JPEG compression, the quality  $q$  is the only variable parameter

$$1011 \hat{I} = JPEG(I, q), \quad (17)$$

1012 which is selected as  $q = q_t^l = q_t^r = q_{t+1}^l = q_{t+1}^r = 6$  for time- and stereo-consistent JPEG  
 1013 compression.

1014 **Elastic Transformation.** The elastic transformation applies an elastic deformation using  
 1015 `torchvision.transforms.v2` with parameters  $\alpha = 110.0$  and  $\sigma = 5.0$  to control the de-  
 1016 formation magnitude and smoothness, while preserving the original frame dimensions.

1017 **Spatter.** The spatter corruption simulates liquid droplets by generating a liquid layer from Gaussian  
 1018 noise, applying blur and thresholding, and blending it with the original image using a predefined  
 1019 color.  
 1020

1021 **Frost.** The frost corruption overlays a frost texture onto the image by randomly selecting and resiz-  
 1022 ing a pre-stored frost image and blending it with the input to create an icy appearance.  
 1023

1024 **Snow and Rain.** The implementation for snow and rain is based on Schmalfuss et al. (2023), with  
 1025 methodological and performance improvements. On the methodological side, we replaced additive  
 1026 blending with order-independent alpha blending (Meshkin’s Method, McGuire & Bavoil (2013))  
 1027 and included global illumination (Halder et al., 2019) in the color rendering. Also, we expanded the

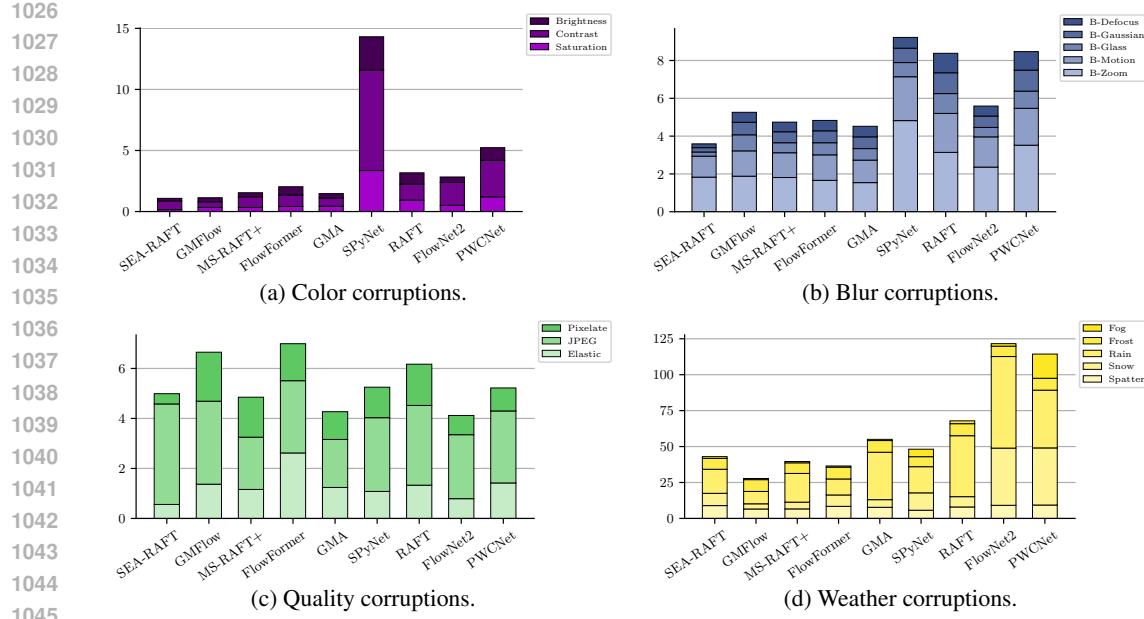


Figure 7: Additional results on accumulated corruption robustness  $R_{\text{EPE}}^c$  for optical flow models over corruption classes color, blur, quality, and weather. More results are in Fig. 4.

monocular two-step motion simulation to multi-step stereo images. On the performance side, we introduce an efficient parallel particle initialization and improve the parallel processing performance.

**Fog.** Fog is modelled using the Koschmieder model from Wiesemann & Jiang (2016) as

$$\hat{I} = I \cdot e^{-\frac{D \cdot \ln(20)}{d_m}} + l \cdot (1 - e^{-\frac{D \cdot \ln(20)}{d_m}}), \quad (18)$$

where  $d_m$  is the visibility range and  $l$  the luminance of the sky. For RobustSpring, we use  $d_m = 45$  and  $l = 0.8$ . As it directly depends on the depth  $D$ , this is depth-consistent and due to its integration into the 3D scene it is also stereo- and time-consistent.

#### A.4 ADDITIONAL EXPERIMENTAL RESULTS

Below, we expand on the experiments in the main paper and provide supplementary results for our major experiments.

##### A.4.1 CORRUPTION ROBUSTNESS BY CORRUPTION GROUP

Figure 7 shows the corruption robustness  $R_{\text{EPE}}^c$  for each optical flow method across the remaining four corruption groups in addition to Fig. 4 in the main paper. It underlines the varying degrees of robustness of the evaluated methods against specific types of corruption.

##### A.4.2 ACCURACY VS. MEDIAN CORRUPTION ROBUSTNESS

In Fig. 8 we show the accuracy-robustness evaluation with the *Median* corruption robustness, to complement Fig. 4c which uses the average corruption robustness. Even though the robustness ranking of methods varies between average and median corruption robustness, *cf.* Sec. 4.2 for a discussion on ranking differences, the general trend that corruption robustness and accuracy are weakly correlated remains. However, there is still no clear winner, and an accuracy-robustness tradeoff persists among particularly accurate or robust methods.

##### A.4.3 SCHULZE PAIRWISE COMPARISON MATRIX

The Schulze method is a ranking algorithm used to determine the most preferred candidate based on pairwise comparisons. We include a pairwise comparison matrix in Table 5 for our ranking.

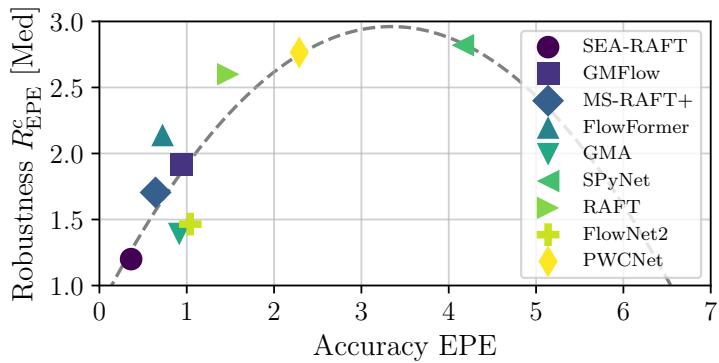


Figure 8: Accuracy vs. robustness of optical flow methods, measured as EPE and median  $R_{\text{EPE}}^c$ . Small values indicate accurate and robust methods. The dashed line represents a quadratic polynomial fit. Fig. 4c shows the average  $R_{\text{EPE}}^c$ .

Table 5: Pairwise comparison matrix for the Schulze method.

	SEA-RAFT	GMFlow	GMA	MS-RAFT+	FlowFormer	SPyNet	RAFT	PWCNet	FlowNet2
SEA-RAFT	0	14	14	14	14	17	17	15	19
GMFlow	6	0	9	14	9	10	16	9	14
MS-RAFT+	6	9	0	15	11	11	19	12	15
FlowFormer	6	6	5	0	3	12	16	8	13
GMA	6	11	9	17	0	12	20	10	15
SPyNet	3	10	9	8	8	0	13	4	13
RAFT	3	4	1	4	0	7	0	4	9
FlowNet2	5	10	8	12	10	16	16	0	18
PWCNet	1	6	5	7	5	7	11	2	0

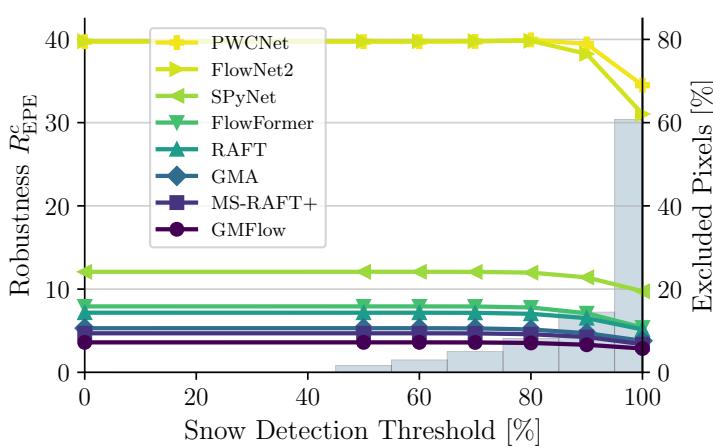


Figure 9: Stability of corruption robustness  $R_{\text{EPE}}^c$  for snow corruption. Analogous rain results in Fig. 5a.

The table shows how often the method in row  $i$  is better than the method in column  $j$ , based on the number of corruptions where method  $i$  achieves a lower error than method  $j$ . The ranking process consists of the following steps:

**1) Constructing the Pairwise Comparison Matrix.** For each pair of methods, we count how many times one method achieves a lower EPE than the other across different corruptions. If method  $A$  has a lower EPE than method  $B$  in a given corruption, the corresponding entry in the matrix is incremented.

**2) Computing the Strongest Paths.** We define the strength of a path from method  $A$  to method  $B$  as the number of cases where  $A$  outperforms  $B$ . The strongest paths between methods are determined by considering indirect paths: if method  $A$  is better than method  $B$ , and method  $B$  is better than method  $C$ , then the strength of the indirect path from  $A$  to  $C$  is considered.

**3) Determining the Final Ranking.** Method  $A$  is ranked higher than method  $B$  if the strongest path from  $A$  to  $B$  is stronger than the strongest path from  $B$  to  $A$ . This ensures that even if a method loses to another in some comparisons, it can still be ranked higher if it consistently performs well against other methods.

#### A.4.4 CORRUPTION ROBUSTNESS ON SNOW

Finally, we complement the evaluation of our corruption robustness metric in the presence of rain in Fig. 5a with the corresponding evaluation in the presence of snow in Fig. 9. The results with rain translate to snow with the following minor differences: Because snow has less motion blur than rain, it covers fewer pixels (60% of all pixels vs. 90% for rain). For snow, the score drops a bit more than for rain when object pixels are excluded ( $\leq 25\%$  drop vs.  $\leq 5\%$  for rain), potentially as a consequence of the increased object opacity for snow particles. Still, the background error ( $\geq 75\%$  contribution to corruption robustness) dominates the score, and the robustness ranking for optical flow methods remains stable, whether snow pixels are included in the score calculation or not. Hence, the additional evaluation on snow further substantiates the stability and expressiveness of corruption robustness as an evaluation metric.

#### A.4.5 REPRODUCIBILITY

We compute the noise in the KITTI (Geiger et al., 2012) frames according to Immerkaer (1996), and select the top 10% noisiest frames. The final frames and their noise levels are listed in Tab. 6.

1188  
 1189 Table 6: Top 10% noisiest KITTI frames with noise estimation from Immerkaer (1996). Frames  
 1190 from kitti15/dataset/training/image\_2.

Frame	000061_10.png	000060_10.png	000065_10.png	000068_10.png	000143_10.png	000069_10.png	000174_10.png	000142_10.png	000107_10.png	000064_10.png	000063_10.png	000066_10.png	000067_10.png	000062_10.png	000055_10.png	000154_10.png	000158_10.png	000157_10.png	000054_10.png	
Noise	3.54	3.22	3.07	2.92	2.80	2.79	2.77	2.74	2.69	2.66	2.62	2.61	2.61	2.60	2.48	2.47	2.46	2.46	2.43	2.41

## 1200 A.5 DEPTH AND EXTRINSICS ESTIMATION FOR CORRUPTIONS

1201  
 1202 **Disparity and Depth.** For depth-dependent corruptions, we estimate disparity maps with MS-  
 1203 RAFT+ (Jahedi et al., 2022; 2024). On the Spring validation set, we find a mean EPE of 1.09 px  
 1204 and D1-all of 5.60%. Mean EPE measures the average Euclidean distance between predicted and  
 1205 ground-truth disparities, while D1-all reports the percentage of pixels with disparity error larger  
 1206 than 3 px or 5% of the ground-truth disparity. These errors are sufficiently small to support realistic  
 1207 augmentations. Note that depth estimates are used by only 3 of the 20 corruption types (fog, snow,  
 1208 and rain; motion blur uses depth only implicitly).

1209  
 1210 **Extrinsics.** We estimate camera extrinsics using COLMAP 3.8. As trajectories are not expressed  
 1211 in a common reference frame, quantitative pose metrics such as absolute trajectory error (ATE) or  
 1212 relative pose error (RPE) cannot be meaningfully computed without a rigid alignment step, which  
 1213 itself introduces error. Extrinsics are required for only 2 corruptions (snow, rain). We therefore  
 1214 evaluate them qualitatively: the resulting augmentations exhibit realistic motion behavior for both  
 1215 effects.

## 1216 A.6 USER STUDY ON SSIM THRESHOLDS

1217  
 1218 The SSIM thresholds of 0.7 for color, blur, quality, and weather corruptions and 0.2 for noise were  
 1219 chosen empirically to yield comparable perceptual corruption strength across categories.

1220  
 1221 We conducted a focused perceptual study to validate the choice of corruption strengths and their as-  
 1222 sociated SSIM values in RobustSpring. For five representative corruptions (one per main corruption  
 1223 family: contrast, fog, Gaussian blur, Gaussian noise, and pixelation), participants repeatedly viewed  
 1224 a clean Spring frame alongside three corrupted versions. See Figs. 10 and 11 for examples. They  
 1225 then selected the corrupted image that looked most like a “reasonably” corrupted version of the clean  
 1226 frame. This means the image is clearly degraded, inducing a distribution shift, yet still recognizably  
 1227 close to the original. Across 14 user studies, with 10 samples each, for a total of 140 comparisons,  
 1228 we observed agreement rates ranging from 70% to 100% from each user for the SSIM values used in  
 1229 our study. The average agreement was 87.14%, which suggests that the SSIM ranges used for these  
 1230 representative corruptions align with human perception of realistic, moderate corruption strengths.

1231  
 1232 In the user study, the web interface displayed one clean reference image from Spring and three  
 1233 corresponding corrupted versions of the same image on each page. These images were obtained  
 1234 by applying a single corruption type at three different intensities. We applied thresholds for three  
 1235 levels of severity: low, medium, and high, with corresponding SSIM values of 0.95, 0.7, and 0.2,  
 1236 respectively. The three corrupted images were randomly ordered, and participants were asked to  
 1237 select one image per page that they perceived as the most reasonably corrupted version of the clean  
 1238 image. The corruption should be clearly visible and sufficient to cause a distribution shift but not so  
 1239 strong that the scene appears implausible or too different from the original image. This setup was  
 1240 repeated for the five corruption types listed above, and for each page and participant, the selected  
 1241 intensity was recorded. This allowed us to compute per-participant agreement with the nominal  
 1242 level and the aggregate agreement statistics reported above. In each scenario, the nominal value  
 1243 corresponds to the SSIM value chosen for the respective corruption in our work.

1242

1243

1244

1245

1246

1247

1248

1249

1250

1251

1252

1253

1254



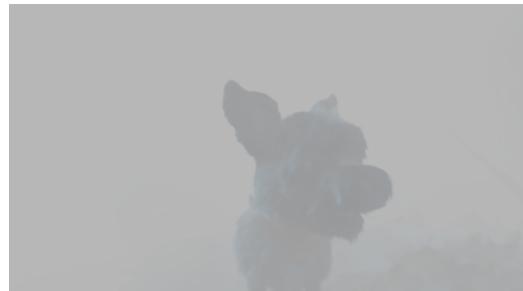
(a) Clean image.



(b) Low severity.



(c) Medium severity (best).



(d) High severity.

1264

1265

1266

1267

1268

1269

1270

1271



(a) Clean image.



(b) Low severity.



(c) Medium severity.



(d) High severity (best).

1292

1293

1294

1295

Figure 10: Clean image and three corrupted versions with different severity levels for fog used in the user study on SSIM thresholds. “Best” is the one we choose in RobustSpring.

1296 **A.7 MOTION RANGE OF SPRING**  
1297

1298 As discussed in the Spring paper (Mehl et al., 2023b), the dataset covers a wider range of motion  
1299 than standard benchmarks. While Sintel (Butler et al., 2012) and KITTI (Menze & Geiger, 2015)  
1300 align with Spring for smaller motions (u:  $-500$  to  $500$ , v:  $-250$  to  $250$ ), Spring extends the range  
1301 up to  $1700$  (u) and  $-750$  (v). This corresponds to motion magnitudes about 1.5–3 times larger than  
1302 those in Sintel and KITTI. We highlight this here as one reason for selecting Spring as the foundation  
1303 of RobustSpring.

1304 **A.8 INTEGRATION INTO THE SPRING BENCHMARK WEBSITE**  
1305

1306 To illustrate how RobustSpring can extend the existing Spring benchmark infrastructure, we provide  
1307 mock-ups of the website with our additional robustness results. These examples demonstrate how  
1308 robustness values can be integrated alongside existing accuracy metrics.  
1309

1310 **Overview Page.** Fig. 12 shows the modified overview page for optical flow methods. Three new  
1311 columns report the average robustness scores. This extension allows users to compare methods  
1312 across both accuracy and robustness at a glance.

1313 **Method Detail Page.** Fig. 13 displays the detail page for a single optical flow method. In addition  
1314 to the accuracy results already present in Spring, robustness values are provided per corruption type,  
1315 as well as aggregated average and median robustness values. This detailed view enables users to  
1316 assess strengths and weaknesses of individual methods with respect to specific corruptions.

1343 Figure 12: Modified Spring benchmark overview page for optical flow. Newly added robustness  
1344 columns are highlighted in red.  
1345  
1346  
1347  
1348  
1349

1350  
1351  
1352  
1353  
1354  
1355  
1356  
1357  
1358  
1359  
1360  
1361  
1362  
1363  
1364  
1365  
1366  
1367  
1368  
1369  
1370  
1371  
1372  
1373  
1374  
1375  
1376  
1377  
1378  
1379  
1380  
1381  
1382  
1383  
1384  
1385  
1386  
1387  
1388  
1389  
1390  
1391  
1392  
1393  
1394  
1395  
1396  
1397  
1398  
1399  
1400  
1401  
1402  
1403



Spring: L. Mehl, J. Schmaljuss, A. Jähedi, Y. Nalivayko, A. Bruhn — University of Stuttgart

Download | Stereo | Optical Flow | Scene Flow | Submit | FAQ

Not logged in | Login | Create Account

### GMFlow

Nov. 1, 2022, 2 p.m. — Public

submitted by spring team | H. Xu, J. Zhang, J. Cai, H. Rezatofighi, and D. Tao. "GMFlow: Learning Optical Flow via Global Matching." In IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 2022.

Code: <https://github.com/haoxiu/gmflow>

#### Optical Flow

<i>1px total</i>	<i>1px low-detail</i>	<i>1px high-detail</i>	<i>1px matched</i>	<i>1px unmatched</i>	<i>1px rigid</i>	<i>1px non-rigid</i>	<i>1px not sky</i>	<i>1px sky</i>	<i>1px s0-10</i>	<i>1px s10-40</i>	<i>1px s40+</i>
10.355	9.935	76.613	9.060	63.949	6.800	37.258	8.952	31.680	5.412	9.901	52.944

<i>EPE total</i>	<i>EPE low-detail</i>	<i>EPE high-detail</i>	<i>EPE matched</i>	<i>EPE unmatched</i>	<i>EPE rigid</i>	<i>EPE non-rigid</i>	<i>EPE not sky</i>	<i>EPE sky</i>	<i>EPE s0-10</i>	<i>EPE s10-40</i>	<i>EPE s40+</i>
0.945	0.872	12.358	0.785	7.577	0.559	3.863	0.941	1.000	0.439	0.728	5.764

<i>Fl total</i>	<i>Fl low-detail</i>	<i>Fl high-detail</i>	<i>Fl matched</i>	<i>Fl unmatched</i>	<i>Fl rigid</i>	<i>Fl non-rigid</i>	<i>Fl not sky</i>	<i>Fl sky</i>	<i>Fl s0-10</i>	<i>Fl s10-40</i>	<i>Fl s40+</i>
2.952	2.695	43.455	2.374	26.853	1.497	13.955	2.753	5.966	0.923	3.804	17.621

<i>WAUC total</i>	<i>WAUC low-detail</i>	<i>WAUC high-detail</i>	<i>WAUC matched</i>	<i>WAUC unmatched</i>	<i>WAUC rigid</i>	<i>WAUC non-rigid</i>	<i>WAUC not sky</i>	<i>WAUC sky</i>	<i>WAUC s0-10</i>	<i>WAUC s10-40</i>	<i>WAUC s40+</i>
82.337	82.657	32.030	83.256	44.305	84.908	62.883	82.983	72.515	86.200	81.393	52.570

#### Robustness

	<i>Corruption</i>	<i>robust EPE</i>	<i>robust 1px</i>	<i>robust Fl</i>
<b>Summary</b>	Average ( $\pm$ std)	2.98 ( $\pm$ 2.70)	40.89 ( $\pm$ 27.91)	14.68 ( $\pm$ 11.91)
	Median	1.92	48.35	13.83
<b>Color</b>	Brightness	0.33	3.31	1.12
	Contrast	0.46	6.71	1.71
	Saturate	0.34	3.30	0.96
<b>Blur</b>	Defocus	0.53	6.17	1.45
	Gaussian	0.66	7.77	1.88
	Glass	0.85	20.87	1.82
<b>Noise</b>	Motion	1.34	18.35	7.51
	Zoom	1.88	35.80	9.90
	Geussian	4.70	57.95	21.67
<b>Quality</b>	Impulse	6.64	66.14	28.70
	Speckle	3.90	62.01	20.64
	Shot	3.52	56.71	17.77
<b>Weather</b>	Pixelate	1.96	68.09	18.71
	JPEG	3.32	83.54	27.92
	Elastic	1.37	40.00	6.89
<b>Fog</b>	Fog	0.80	14.42	5.32
	Frost	8.20	63.96	29.96
	Rain	8.60	64.20	32.72
<b>Spatter</b>	Snow	3.60	70.60	29.90
	Spatter	6.58	67.90	27.09

Lukas Mehl, Jenny Schmaljuss, Azin Jähedi, Yaroslava Nalivayko, Andrej Bruhn:  
"Spring: A High-Resolution High-Detail Dataset and Benchmark for Scene Flow, Optical Flow and Stereo".  
In IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 2023.

 University of Stuttgart  
Institute for Visualization and Interactive Systems

 SFB-TRR 161  Institute for Visual Computing

The Spring movie assets by Blender Foundation are licensed under CC BY 4.0. The authors thank Andy Golatczky, Francesco Sedita and the whole Blender Studio team. Lukas Mehl acknowledges funding by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – Project-ID 251604679 – 1881133 Quantitative Methods for Visual Computing (QVC). Jenny Schmaljuss and Azin Jähedi acknowledge support from the International Max-Planck Research School for Intelligent Systems (IMPRS-IS).

[Contact](#) | [Imprint](#) | [Privacy](#)

Figure 13: Modified Spring benchmark method detail page for optical flow. Robustness values are reported for each corruption, as well as aggregated average and median scores. Newly added sections are highlighted in red.