

# DEFLOW: Self-supervised 3D Motion Estimation of Debris Flow

Liyuan Zhu<sup>1</sup> Yuru Jia<sup>1</sup> Shengyu Huang<sup>1</sup> Nicholas Meyer<sup>1</sup>  
Andreas Wieser<sup>1</sup> Konrad Schindler<sup>1</sup> Jordan Aaron<sup>2,3</sup>

<sup>1</sup>Institute of Geodesy and Photogrammetry, ETH Zurich

<sup>2</sup>Swiss Federal Institute for Forest, Snow and Landscape Research WSL

<sup>3</sup>Geological Institute, ETH Zurich

## Abstract

Existing work on scene flow estimation focuses on autonomous driving and mobile robotics, while automated solutions are lacking for motion in nature, such as that exhibited by debris flows. We propose DEFLOW, a model for 3D motion estimation of debris flows, together with a newly captured dataset. We adopt a novel multi-level sensor fusion architecture and self-supervision to incorporate the inductive biases of the scene. We further adopt a multi-frame temporal processing module to enable flow speed estimation over time. Our model achieves state-of-the-art optical flow and depth estimation on our dataset, and fully automates the motion estimation for debris flows. Source code and dataset are available at [project page](#).

## 1. Introduction

Cameras and LiDAR sensors are complementary for 3D scene understanding, consequently many algorithms designed for autonomous driving and mobile robotics use these two modalities. However, little attention has been paid to natural scenes. There is potential to use these sensors to analyse natural processes, and debris flows are particularly amenable to this [1, 16, 22]. Debris flows are extremely rapid flows of soil, water and woody debris that can reach velocities in excess of 5 m/s, and travel for multiple kilometers [18]. They are a considerable hazard in mountainous regions and cause costly disasters every year [34]. Reducing the impact of debris-flows requires a better understanding of the fundamental mechanisms which drive motion, and progress has been previously limited by a lack of high temporal and spatial resolution observations of debris-flow velocity. In the present work, we analyse a debris-flow dataset captured with a camera-LiDAR setup (see Fig. 1) and estimate dense surface velocity fields.

Two broad categories of methods can be considered when analyzing natural debris-flows using camera and Li-

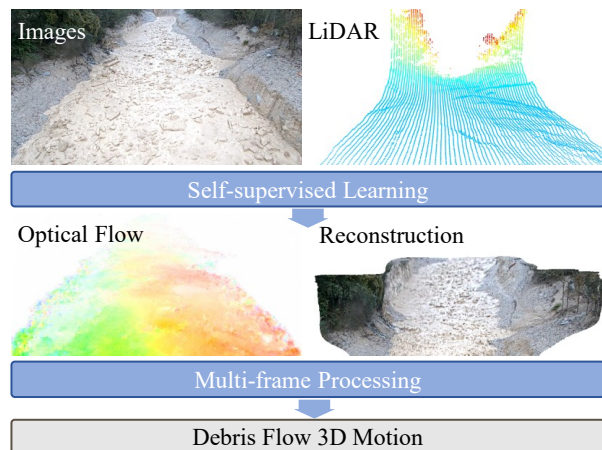


Figure 1. **Debris flow estimation from images and sparse LiDAR points.** Our self-supervised method automates the estimation of dense scene depth and flow. Additionally, a multi-frame temporal processing module recovers flow velocity profiles over time.

DAR sensors: flow estimation techniques or 3D scene flow. Recent work on fluid flow estimation [26, 27, 55] has focused on embedding physical constraints into Particle Image Velocity (PIV) solutions. The fluid datasets used are often either generated by numerical simulations, or collected under well-constrained laboratory conditions, which makes it difficult to generalize to real-world scenarios. 3D scene flow estimation relies heavily on the rigid body motion prior [5, 15, 17, 44], and point-wise correspondence [15, 17, 28]. These methods are not well-suited to debris flows, as the bulk motion is non-rigid, although locally rigid motion occurs when boulders and woody debris are present on the flow surface [2]. From the above observations, we find that existing methods are hampered by (i) inapplicable physical constraints (ii) over-reliance on rigid motion priors and (iii) an inability to generalize to our hard-to-annotate debris-flow dataset.

To overcome the above shortcomings, we propose a method that is tailored to debris-flow monitoring. With no

direct supervision from manual annotation, we introduce a self-supervised optical flow branch [19, 39] into our network, which exploits the informative and easily trackable pixel features in optical images. Since optical flow estimations only provide 2D pixel-wise motion, we incorporate a depth estimation branch to densify sparse depth observations from LiDAR and lift the optical flow to 3D scene flow. In contrast to [15, 17], we take into account the unique properties of debris flows and of the monitoring system, namely (i) the absence of ego-motion and (ii) the presence of a continuous fluid surface without motion-induced occlusion. Through experimentation, we show that the optical flow task and the depth estimation task mutually reinforce each other and that our method’s performance is further improved by the inclusion of suitable inductive biases and warping-based multi-frame temporal smoothing. Additionally, our method’s lightweight nature makes it easy to train from scratch to analyze new debris flow events. We apply our method to derive dense flow velocity fields of a debris-flow event, at the highest spatial resolution obtained so far for a natural debris flow. Finally, we release our code and debris-flow dataset to the community, as a benchmark for understanding motion in nature.

## 2. Related work

**Scene flow from RGB images.** Recent studies have explored a variety of methods to estimate depth and motion from RGB images. Brickwedde *et al.* [6] propose a probabilistic depth estimation network and compute the motion of the scene by combining multi-view geometry and single-view depth information. Yang *et al.* [53] introduce optical expansion and propose a network that learns to expand optical flow and estimate motion in depth. Hur *et al.* [19] jointly estimate depth and 3D scene flow from monocular images by adopting self-supervised learning.

However, [6, 43] heavily rely on a rigid motion prior for moving vehicles, which does not apply to debris flows. [19, 53] suffer from overfitting to specific camera intrinsics. Their networks are trained on KITTI [11] with known camera intrinsics, but encounter problems when applied to other datasets, *e.g.*, nuScenes [8] or our debris-flow data.

**Camera-LiDAR fusion.** An intuitive way to disambiguate depth scaling is to incorporate direct 3D measurements into image-based pipelines. Many works have thus focused on sensor fusion for scene flow estimation. LiDAR-Flow [4] fuses the LiDAR with a stereo camera to boost the cost computation in the matching process for both stereo and optical flow. Subsequent work like DeepLiDARFlow [36] encapsulates the solution into an end-to-end deep neural network that regresses scene flow from consecutive LiDAR scans and monocular images. CamLiFlow [29] mitigates the sen-

sor fusion problem with a multi-stage fusion framework and estimates point-wise motion for sparse LiDAR point clouds.

**Multi-frame flow processing.** Two-frame methods [6, 19, 29, 36, 53] do not account for the temporal continuity across multiple consecutive frames. In contrast, [42] enforces consistency for both spatial and temporal neighbors in a sliding temporal window, [40] leverages multi-frame consistency to detect moving object regions, and [20] introduces a convolutional LSTM on top of [19] to propagate the hidden states via forward warping, but requires stereo images as supervision. Huang *et al.* [17] align multi-frame 3D scans in a common reference frame, accumulate 3D points on individual objects, and obtain a better decomposition of the dynamic scene. Wang *et al.* [45] encode the spatial-temporal information into a neural trajectory prior. These previous works exploit the geometric layout and rigidity of autonomous driving scenes, but do not generalize well to debris flows. Our work takes temporal information into account and estimates the flow speed across the entire event by warping-based temporal smoothing, which is a stable and efficient choice for debris flows in the absence of strong supervision signals, *e.g.*, stereo images or ground truth scene flow.

**Debris flow estimation.** Debris-flow hazard is largely governed by flow depths and velocities. These parameters have only been measured in a few locations worldwide, often using measurement techniques that have poor spatial and/or temporal resolution [21]. Unlike these previous studies, Aaron *et al.* [1, 2, 37] conducted field measurements to directly measure in-situ parameters of debris flows at high spatial and temporal resolution with time-lapse LiDARs and cameras. Spielmann *et al.* [37] manually measure the front velocity and track individual features such as large boulders and woody debris in the point cloud data. Aaron *et al.* [1] start from images and explore both pixel-level motion and object-level motion: the former method applies particle image velocimetry (PIV) to the video frames, while the latter focuses on tracking boulders and woody debris using CNN-based object detection. Subsequently, both methods project the 2D displacements into the point clouds. Although both LiDAR and camera sensors are employed, the connection between them is not fully exploited with a simple projection. Further, the velocity fields obtained are relatively sparse when compared to what could be obtained with dense optical flow techniques. To this end, we develop tailored sensor-fusion methods to fully utilize the multi-sensor setup for the debris-flow dataset, resulting in much denser flow velocity fields.

## 3. Method

The workflow of our method is illustrated in Fig. 2. Our network starts with two feature encoders, followed by fea-




Figure 2. **Architecture overview.** The inputs of our network are two consecutive camera images and the corresponding (synchronous) range maps generated by a LiDAR sensor. The optical and depth information are encoded by image and sparse depth encoders. Next, a multi-level sensor fusion scheme combines depth features, image features, and the feature correlation volume. The aggregated feature maps are fed into a multi-task decoder to output the depth and optical flow estimates. After the learnable part, we employ deterministic geometric relations to back-project the depth image into the point cloud and compute 3D motion from correspondences defined by the optical flow.

ture warping, correlation volume computation, and multi-level fusion, and finally decodes the resulting latent features to optical flow and depth estimates. § 3.2 and § 3.3 elaborate two necessary losses to supervise the learning of optical flow and depth without ground truth. In § 3.4 and § 3.5, we explain how we leverage two inductive biases specific to debris flow. To finally obtain the 3D motion, post-processing (§ 3.7) is used to convert the optical flow and depth to 3D flow. The flow speed profile over time is estimated after multi-frame temporal smoothing (§ 3.6). We train the model end-to-end with a loss  $\mathcal{L}$  composed of four terms:

$$\mathcal{L} = \lambda_{\text{flow}} \mathcal{L}_{\text{flow}} + \lambda_{\text{depth}} \mathcal{L}_{\text{depth}} + \mathcal{L}_{\text{static}} + \mathcal{L}_{\text{cycle}}. \quad (1)$$

**Problem formulation.** Consider an image sequence  $\mathcal{I} =$

$\{\mathbf{I}^t\}_{t=1}^T$  and a point cloud sequence  $\mathcal{X} = \{\mathbf{X}^t\}_{t=1}^T$  with constant interval in  $T$  epochs. Our objective is to estimate the 3D coordinates  $\mathbf{P}^t = (x, y, z)$  and the flow vector  $\mathbf{V}^t = (\Delta x, \Delta y, \Delta z)$  for every pixel  $\mathbf{p}^t = (u, v) \in \mathbf{I}^t$ . Additionally, the flow velocity over time needs to be accumulated to generate the flow speed profile  $\mathbf{S} = \{\|\mathbf{V}^t\|_2\}_{t=1}^T$ .

### 3.1. Backbone network

We follow [19, 39, 50] and iteratively estimate the optical flow in a coarse-to-fine manner. Our network takes both images and point clouds as inputs, and fuses their features at multiple levels.

**Feature encoders.** The network starts with an image encoder and a depth encoder. Image features  $\mathbf{F}_{\text{image}}^t$  and  $\mathbf{F}_{\text{image}}^{t+1}$  are extracted by a shared image encoder. The depth encoder learns to extract and aggregate the sparse range features  $\mathbf{F}_{\text{depth}}^t$ . Both encoders have a pyramid structure with five layers to preserve global context and local details. The feature pyramid channels are [32, 64, 96, 128, 192] and [8, 16, 24, 32, 64] for image encoder and depth encoder, respectively. At each layer a stride of 2 halves the size of the feature map.

**Multi-level fusion.** We carry out feature fusion at each level of the feature pyramid. The first iteration initializes optical flow  $\mathbf{O}^t$  and depth  $\mathbf{D}^t$  to *zero*. Based on the current optical flow estimates, the feature map  $\mathbf{F}_{\text{image}}^{t+1}$  is warped [47] back to  $t$  to obtain  $\tilde{\mathbf{F}}_{\text{image}}^t$ . Then we compute the correlation volume  $\mathbf{F}_{\text{corr}}^t$  by stacking the feature-wise correlation, between  $\mathbf{F}_{\text{image}}^t$  and  $\tilde{\mathbf{F}}_{\text{image}}^t$  with window size  $l$ . The feature-wise correlation  $\mathbf{f}_{\text{corr}}$  between  $\mathbf{f}_i^t \in \mathbf{F}_{\text{image}}^t$  and  $\tilde{\mathbf{f}}_j^t \in \tilde{\mathbf{F}}_{\text{image}}^t$  is computed as

$$\mathbf{f}_{\text{corr}} = \frac{1}{C} \sum_{m \in [0, C)} [\mathbf{f}_i^t \circ \tilde{\mathbf{f}}_j^t]_m, \quad (2)$$

where  $\circ$  denotes the element-wise product, indices  $i, j$  refer to the  $j^{\text{th}}$  correlation of the  $i^{\text{th}}$  feature, and  $C$  is the channel depth of the features. For a feature map  $\mathbf{F}_{\text{image}}^t \in$

$\mathbb{R}^{W \times H \times N}$ , the correlation volume  $\mathbf{F}_{\text{corr}}^t \in \mathbb{R}^{W \times H \times l^2}$  is generated.  $W$ ,  $H$  and  $N$  correspond to the width, height, and number of channels of the feature map, respectively. Then all the latent features are concatenated,

$$\mathbf{F}_{\text{fuse}}^t = \text{cat}(\mathbf{F}_{\text{image}}^t, \mathbf{F}_{\text{image}}^{t+1}, \mathbf{F}_{\text{depth}}^t, \mathbf{F}_{\text{corr}}^t, \mathbf{O}^t, \mathbf{D}^t), \quad (3)$$

and the fused feature map  $\mathbf{F}_{\text{fuse}}^t$  is fed into the decoders.

**Decoder design.** We adopt separate decoders for depth and flow as described in [20], for better convergence. The fused feature  $\mathbf{F}_{\text{fuse}}^t$  first goes through a shared convolutional layer to output a latent feature map, which is then passed into two separate decoders. The context network from [9, 54] is adopted after each decoder, which outputs the depth and flow estimates.

### 3.2. Optical flow estimation

The optical flow decoder consists of three convolutional layers and a context network (dilated convolutions). Unlike vehicles in KITTI [12] or NuScenes [8], debris-flow motion is non-rigid. It is hard to model debris-flow motion through the Navier-Stokes equation [10] due to significant unknowns related to the underlying mechanisms that control flow properties. Moreover, manual annotation of optical flow on such a dataset is almost infeasible. Therefore, we use an unsupervised loss [3, 23] based on photometric consistency and local smoothness. Given two images ( $\mathbf{I}^t, \mathbf{I}^{t+1}$ ) and an optical flow  $\mathbf{O}_f^t$ , we warp  $\mathbf{I}^{t+1}$  back to  $t$ :

$$\tilde{\mathbf{I}}^t = \omega(\mathbf{I}^{t+1}, \mathbf{O}_f^t), \quad (4)$$

where  $\omega(\cdot, \cdot)$  represents the backward warping function with the first argument being the scalar field to be warped and the second argument the warping correspondences. Under the assumption that the appearance of the same object remains constant, the photometric loss  $\mathcal{L}_{\text{photo}}$  penalizes the photometric differences between  $\mathbf{I}^t$  and  $\tilde{\mathbf{I}}^t$ :

$$\mathcal{L}_{\text{photo}} = \text{SSIM}(\phi(\mathbf{I}^t), \phi(\tilde{\mathbf{I}}^t)), \quad (5)$$

where  $\phi(\cdot)$  is an average pooling module and  $\text{SSIM}(\cdot, \cdot)$  is the structural similarity index [46].

Besides, a first-order regularizer [3]  $\mathcal{L}_{\text{smooth}}$  encourages edge-aware smoothness [19, 31, 49] of the estimates by penalizing the sum of first-order derivatives, weighted by the image gradients:

$$\mathcal{L}_{\text{smooth}} = \frac{1}{H \times W \times C} \sum_{\mathbf{p}} \sum_{i \in \mathbf{I}^t} \|\nabla_i^1 \mathbf{O}^t(\mathbf{p})\|_1 \cdot e^{-\beta \|\nabla_i \mathbf{I}^t(\mathbf{p})\|_1}, \quad (6)$$

with  $\beta = 10$  and  $H \times W \times C$  the image dimensions. The total flow loss is  $\mathcal{L}_{\text{flow}} = \mathcal{L}_{\text{photo}} + \lambda \mathcal{L}_{\text{smooth}}$ , with  $\lambda = 0.15$ . In the forward propagation, we additionally swap the two input images [19, 39] and compute the backward

flow  $\mathbf{O}_b^t$  and its corresponding photometric loss. During training  $\mathcal{L}_{\text{flow}}$  includes the bidirectional flow loss, whereas only forward flow is computed at inference time. In this way we double the training samples, moreover the bidirectional flows are vital to leverage fluid continuity (§ 3.5).

### 3.3. Depth estimation




Figure 3. **A simple and effective way of supervising depth.** A downsampled point cloud serves as input, the full point cloud as the target.

Previous work on unsupervised and zero-shot monocular depth estimation [13, 14, 19, 35] and depth completion [48] does not rely on ground truth (direct depth) as supervision, but instead uses disparity maps acquired from stereo images. In the debris flow dataset, we have neither ground truth nor good stereo pairs. Therefore, we propose a novel way to supervise the depth estimation as shown in Fig. 3. We start from randomly downsampling the input point cloud  $\mathbf{X}^t$  into  $\mathbf{X}_{\text{down}}^t = [\mathbf{x}_i^t, \dots, \mathbf{x}_j^t, \dots, \mathbf{x}_{n_t}^t] \in \mathbb{R}^{3 \times (\eta \times n_t)}$  with  $\eta$  being the downsampling ratio. During training, only the sampled point cloud  $\mathbf{X}_{\text{down}}^t$  is converted to a range map  $\mathbf{R}_{\text{down}}^t$ , and the network makes predictions based on that downsampled map. During back-propagation, the original input point cloud  $\mathbf{X}^t$  serves as ground truth. After converting  $\mathbf{X}^t$  into a sparse range map  $\mathbf{R}^t$ , we bilinearly sample the points with a valid LiDAR footprint from the dense prediction  $\mathbf{D}^t$  and compute the loss from those samples and the sparse range map:

$$\mathcal{L}_1 = \frac{1}{H \times W \times C} \sum_i \|\mathbf{R}^t(i) - \xi(\Theta_{\text{depth}}(\mathbf{I}^t, \mathbf{I}^{t+1}, \mathbf{R}_{\text{down}}^t), \mathbf{R}^t(i))\|_1, \quad (7)$$

where  $\Theta_{\text{depth}}$  represents the depth branch of our network and  $\mathbf{R}_{\text{down}}^t$  is the downsampled range map. The bilinear resampling function  $\xi(\mathbf{X}, y)$  retrieves  $x$  from  $\mathbf{X}$  at the position of  $y$ . As in § 3.2, smoothness regularization  $\mathcal{L}_{\text{smooth}}$  is used in the depth prediction. The depth estimation branch is supervised by the weighted sum  $\mathcal{L}_{\text{depth}} = \mathcal{L}_1 + \lambda \mathcal{L}_{\text{smooth}}$ , with  $\lambda = 0.1$ . The rationale behind this supervision strategy is that (i) the sampled LiDAR points teach the network to preserve the ground truth input values, while (ii) remaining LiDAR points provide a supervision signal for the densification/interpolation from sparse range input to dense output.

### 3.4. Motion segmentation

The debris flow dataset is captured by a static monitoring system. Different from many scene flow methods [15, 17] which subtract the ego-motion before analyzing the dynamic scene, we leverage the prior of *zero* ego-motion, via a loss  $\mathcal{L}_{\text{static}}$ . As described in § 3.2, we compute the forward flow

$$(\mathbf{O}_f^t, \mathbf{D}^t) = \Theta(\mathbf{I}^t, \mathbf{I}^{t+1}, \mathbf{R}_{\text{down}}^t), \quad (8)$$

and the backward flow by swapping the input order

$$(\mathbf{O}_b^t, \mathbf{D}^{t+1}) = \Theta(\mathbf{I}^{t+1}, \mathbf{I}^t, \mathbf{R}_{\text{down}}^{t+1}), \quad (9)$$

where  $\Theta$  are the parameters of the network. With bidirectional flows  $\mathbf{O}_f^t$  and  $\mathbf{O}_b^t$  and depth estimates  $\mathbf{D}^t$  and  $\mathbf{D}^{t+1}$ , we define a prior that favours a static sensor system. After distinguishing static pixels from moving ones with a simple threshold  $\varepsilon$ , the loss demands that depth estimates in the static part should be consistent, and penalizes deviations from that soft constraint:

$$\mathcal{L}_{\text{static}} = \frac{\sum_{\mathbf{p}} (\mathcal{M}^t(\mathbf{p}) \cdot \|\mathbf{D}^{t+1}(\mathbf{p}) - \mathbf{D}^t(\mathbf{p})\|_1)}{\sum_{\mathbf{q}} \mathcal{M}^t(\mathbf{q})}. \quad (10)$$

Here  $\mathbf{p}$ ,  $\mathbf{q}$  represents pixel coordinates in the image, and  $\mathcal{M}^t$  denotes a binary mask that is 1 for static pixels and 0 for moving ones. The loss  $\mathcal{L}_{\text{static}}$  builds a bridge so that optical flow estimation and depth estimation can mutually supervise each other. Inaccurate optical flow estimates lead to wrong motion segmentation, and consequently,  $\mathcal{L}_{\text{static}}$  increases due to the misalignment of non-static pixels. Unstable or inaccurate depth estimates in static regions also cause the increase of  $\mathcal{L}_{\text{static}}$ .

### 3.5. Fluid continuity

One long-standing problem in optical flow estimation is occlusion. In this work, we use the inconsistency [31, 51, 52] of forward-backward optical flow to achieve an inductive bias towards fluid continuity. The sensor system is static, and the motion only occurs on the fluid surface with approximately in-plane motion. The continuity of the fluid surface results in almost no occlusion, so we integrate this prior into the loss  $\mathcal{L}_{\text{cycle}}$  as

$$\mathcal{L}_{\text{cycle}} = \frac{1}{H \times W \times C} \sum_{\mathbf{p}} \|\mathbf{O}_f^t(\mathbf{p}) + \omega(\mathbf{O}_b^t, \mathbf{O}_f^t)(\mathbf{p})\|_1, \quad (11)$$

with  $\omega(\cdot, \cdot)$  the backward warping function from § 3.2.

### 3.6. Multi-frame temporal smoothing

Debris flow motion analysis, *e.g.*, calculating flow speed profiles over time, requires more than two-frames. We add a post-processing module to achieve a temporally smoother




Figure 4. Schematic illustration of temporal smoothing for flow speed estimation.

flow speed estimate, as shown in Fig. 4. Given consecutive images  $\{\mathbf{I}^t\}_{t=0}^3$ , and their bidirectional flow estimates  $\{\mathbf{O}_n^m\}_{n=f,b}^{m=0,1,2}$ , our goal is to smooth the flow  $\mathbf{O}_f^1$ . The assumption is that the velocity of object points remains constant over a short time interval. We start with warping  $\mathbf{O}_f^0$  forward to the next frame  $\mathbf{O}_f^{1'} = \omega(\mathbf{O}_f^0, \mathbf{O}_b^0)$ . Then  $\mathbf{O}_f^2$  is warped backward  $\mathbf{O}_f^{1''} = \omega(\mathbf{O}_f^2, \mathbf{O}_b^1)$ . We compute the smoothed  $\tilde{\mathbf{O}}_f^1$  as

$$\tilde{\mathbf{O}}_f^1 = \lambda_0 \mathbf{O}_f^{1'} + \lambda_1 \mathbf{O}_f^1 + \lambda_2 \mathbf{O}_f^{1''}, \quad (12)$$

where  $\lambda_0 + \lambda_1 + \lambda_2 = 1$ . In the end, we apply temporal smoothing to the per-frame estimates to generate the flow speed profile.

### 3.7. Lifting optical flow to 3D

As a final step we transform optical flow to 3D scene flow with the help of depth. The camera matrix  $\mathbf{K}_{3 \times 3}$  and the camera-LiDAR transformation matrix  $\mathbf{T}_{3 \times 4} : [\mathbf{R} \mid \mathbf{t}]$  are known from calibration. Given a depth estimate  $z_i^t$ , every pixel  $\mathbf{p}_i^t \in \mathbf{I}^t$  is back-projected into a 3D point  $\mathbf{P}_i^t$  as

$$\mathbf{P}_i^t = z_i^t \mathbf{R}^T \mathbf{K}^{-1} \mathbf{p}_i^t - \mathbf{R}^T \mathbf{t}. \quad (13)$$

The same transformation is applied to the second frame  $\mathbf{I}^{t+1}$ . Consider two point clouds  $\{\mathbf{P}_i^t\}_{i=0}^{w \times h}$  and  $\{\mathbf{P}_i^{t+1}\}_{i=0}^{w \times h}$ , we determine the motion for every point in frame  $t$  by subtracting its positions in epoch  $t+1$  and  $t$  as in Eq (2).

### 3.8. Implementation details

We have implemented our network in PyTorch [33] and trained it from scratch on a single RTX 6000/3090 24G GPU. During training the loss of Eq (1) is minimized with the Adam optimizer [24], with hyper-parameter  $\beta_1 = 0.9$  and an initial learning rate of 0.0004. We train the network for 30 epochs with a batch size of 8, the learning rate decays by half after 5, 7 and 15 epochs. To enlarge the training set, we adopt geometric data augmentation, including random cropping and flipping for images and range maps simultaneously. The original image size is  $1920 \times 1080$ , which we crop it to  $1600 \times 960$ , making sure the dimensions are divisible by  $2^5 = 32$ .

## 4. Experiments

We first introduce the debris flow dataset and our evaluation metrics for optical flow and depth estimation. Then

we compare our model to several baselines. Finally, we perform an ablation study to validate our network architecture and loss design.

#### 4.1. Debris flow dataset

The debris flow event was captured by a camera-LiDAR setup at Illgraben, Switzerland [30]. The whole event lasted for approximately 30 minutes and consisted of three stages. Stage (i) features a watery pre-surge, stage (ii) marks the arrival of the flow front with boulders and woody debris and is followed by stage (iii) with a finer-grained slurry. This event is described and analysed in detail by [1, 2, 37]. The point clouds are captured by a 64-beam Ouster OS1-64 (Gen-1) LiDAR at a frequency of 10 Hz with each scanline consisting of 2048 points. Along with the LiDAR, two cameras are also deployed, capturing high-resolution videos of the event at 25 Hz. The rigid body transformation between the two sensors as well as the camera intrinsics are known from calibration. The dataset consists of a total of 6000 ordered frames (*i.e.*, the first 10 minutes of the flow). We divide it into chunks of 60 frames, of which the first 45 are assigned to the training set and the last 15 are used for evaluation. This split avoids excessive (temporal) correlation between the training and test sets, but still allows one to evaluate model performance at different stages of the debris flow event. As the two cameras are poorly synchronized and do not form a good stereo baseline, we build and evaluate our method only with the top-view camera and LiDAR.

#### 4.2. Evaluation setting

**Optical flow metrics.** As our dataset lacks a direct ground truth for optical flow, typical metrics such as angular error (AE) [3] or end-point error (EPE) [32] are not applicable; while  $\mathcal{L}_{\text{photo}}$  is not an independent metric, since it is part of the loss minimized during training. We resort to the frame interpolation metric of Baker *et al.* [3]. With the optical flow estimate  $\mathbf{O}_f^t$ , we reconstruct the first frame  $\tilde{\mathbf{I}}^t$  by back-warping the second frame  $\mathbf{I}^{t+1}$ . The root-mean-square deviation (RMSD) between the original image and the reconstructed image is computed as

$$\text{RMSD} = \sqrt{\frac{1}{H \times W \times C} \sum_{\mathbf{p}} \left| \mathbf{I}^t(\mathbf{p}) - \tilde{\mathbf{I}}^t(\mathbf{p}) \right|_2^2}. \quad (14)$$

Furthermore, we also utilize the ternary census transform loss [31, 38, 56] in our experimental analysis, as it is robust against changes in illumination.

**Depth estimation metrics.** For the evaluation of depth estimates, we use the depths from LiDAR points as ground truth and compute the mean absolute error within 10 m, 30 m, and 50 m distance from the LiDAR sensor ( $\text{MAE}_{10}$ ,

$\text{MAE}_{30}$  and  $\text{MAE}_{50}$ ) as well as the absolute relative error (Abs. Rel.).

**Baselines.** We establish four baselines for comparison purposes. For optical flow estimation, we employ RAFT [41], a leading supervised method, and compare its pretrained and finetuned model to ours. We also include a comparison with an unsupervised optical flow network [31] as another baseline. To evaluate depth estimation, we construct an image-only baseline, by removing the depth encoder branch in our network. We also use IP-Basic [25] as another depth completion baseline. Lastly, we establish the final baseline by comparing our estimates to the PIV-based flow speed estimates of Aaron *et al.* [1] for the same dataset.

#### 4.3. Main results

Method	RMSD ↓	Census Loss ↓
UnFlow [31]	19.58	0.146
RAFT-Sintel [41]	11.75	0.156
RAFT-Sintel-ft [41]	10.18	0.147
DeFlow-Cam ( <i>Ours</i> )	<u>7.25</u>	<u>0.109</u>
DeFlow-Fusion ( <i>Ours</i> )	<b>7.02</b>	<b>0.106</b>

Table 1. Results of optical flow estimation.

**Optical flow estimation.** We test unsupervised UnFlow [31], which shows weak performance on the debris flow dataset. We also evaluate a state-of-the-art optical flow method, RAFT [41], on the debris flow dataset. We first evaluate the model pretrained on Sintel [7]. Furthermore, we also finetuned that model, using the same flow loss  $\mathcal{L}_{\text{flow}}$  and the same hyperparameter settings as in our new method, to conduct a fair comparison. After finetuning, we can see a performance gain of RAFT on the debris flow dataset (Tab. 1), which confirms a domain gap between canonical flow datasets and the newly available debris flow. Both our hard baseline DeFlow-Cam and full model DeFlow-Fusion outperform the finetuned RAFT model by a large margin of  $\approx 30\%$ , see Tab. 1. In terms of network design, our architecture internally downsamples the raw input by a factor  $\times 2$ , while RAFT downsamples to  $\times 8$  the original size. Both RAFT and our network adopt an iterative refinement scheme. We believe that the salient and variable local texture information in debris flow is better preserved by our model due to less downsampling and consequently higher-resolution feature maps.

**Depth estimation.** We initiated our study with a camera-only baseline similar to [19] without the top depth encoder in Fig. 2 to assess whether a single camera is capable of providing sufficient information for the task. In Tab. 2, we denote the camera-only model and the camera-LiDAR fusion as *DeFlow-Cam* and *DeFlow-Fusion*, respectively. To

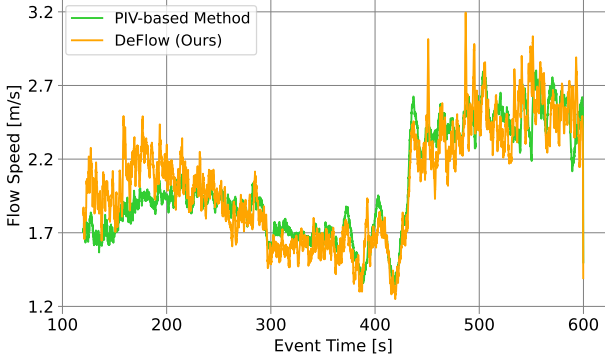


Figure 5. **Flow speed estimates over time.** Comparison between our method (DeFlow) and PIV method of Aaron *et al.* [1, 2].

ensure a fair comparison, we trained both models using only the two essential loss terms  $\mathcal{L}_{\text{flow}}$  and  $\mathcal{L}_{\text{depth}}$ . The DeFlow-Fusion model achieves sub-decimeter accuracy, surpassing the performance of DeFlow-Cam significantly. By fusing the depth branch into our network, the Abs. Rel is reduced  $6\times$ , from 6.3% to 0.9%. DeFlow-Fusion also outperforms the depth completion baseline IP-Basic [25] by approximately 30% in all metrics. Thus, we conclude that a camera-LiDAR fusion approach is essential for debris flow estimation, where the sparse LiDAR input provides strong guidance for depth estimation. As shown in Tab. 1, the integration of LiDAR not only improves depth estimation, but also enhances the optical flow. We believe this performance gain is due to the LiDAR branch and sensor fusion module providing more information about the 3D scene structure, which allows the network to better understand debris-flow motion.

Method	MAE <sub>10</sub> [m] ↓	MAE <sub>30</sub> [m] ↓	MAE <sub>50</sub> [m] ↓	Abs. Rel.[%] ↓
IP-Basic [25]	0.06	0.17	0.22	1.1
DeFlow-Cam	0.38	0.83	1.02	6.3
DeFlow-Fusion	<b>0.04</b>	<b>0.10</b>	<b>0.14</b>	<b>0.6</b>

Table 2. **Depth estimation on our debris flow dataset.**

**Flow speed estimation.** Our approach determines the pixel-wise 3D motion for every frame with multi-frame temporal smoothing. It can therefore provide velocity and speed estimates for any region in the flow channel. Here we compute an estimate of the flow speed within a vertical bounding box  $\{(x, y, z) \mid x \in [-1.0, 1.0], y \in [19.0, 21.0]\}$  in the LiDAR frame, the same region also used in [1]. The average flow speed in the bounding box during the whole event is computed and compared to the results of Aaron *et al.* [1], derived with a PIV-based method and validated with manual feature measurements (Fig. 5). The two results are similar in terms of the overall trend, the absolute velocities, as well as the timing of peaks.




Figure 6. **Spatial distribution of flow speeds across the channel.**

**Channel speed distribution.** To better understand how flow velocity is distributed in space, we investigate the flow speed along a horizontal cross-section of the channel. Specifically, we plot the velocity against the x-axis with  $y \in [470, 490]$  at three different times, represented by epoch 1200, 3000, and 4800 (Fig. 6). During the first two epochs (1200 and 3000), we observe a nearly constant cross-channel velocity profile, whereas during the final epoch the shape of the velocity profile changes. This corresponds to a large rise in overall flow velocity (Fig. 5), and may reflect changes in flow composition [2].

**Qualitative results.** We show qualitative examples of optical flow fields and scene reconstructions generated by our approach. Our model successfully segments the motion of large features, such as boulders, in the flow field (Fig. 7). It further distinguishes between dynamic and stationary regions, and reconstructs the 3D scene geometry, as shown in Fig. 7. These results represent the densest flow field obtained so far for a natural debris flow, and further processing will make it possible to better understand the interaction of large particles with the surrounding fine-grained slurry.

#### 4.4. Ablation Study

Context	Modules		Optical Flow ↓	Depth Estimation ↓			
	$\mathcal{L}_{\text{static}}$	$\mathcal{L}_{\text{cycle}}$	RMSD	MAE <sub>10</sub> [m]	MAE <sub>30</sub> [m]	MAE <sub>50</sub> [m]	Abs. Rel.[%]
<b>x</b>	<b>x</b>	<b>x</b>	7.018	0.053	0.139	0.196	0.83
<b>✓</b>	<b>x</b>	<b>x</b>	6.969	<u>0.047</u>	0.155	0.207	0.90
<b>✓</b>	<b>x</b>	<b>✓</b>	7.024	0.051	0.132	<u>0.179</u>	0.82
<b>✓</b>	<b>✓</b>	<b>x</b>	<u>6.871</u>	0.049	0.137	0.196	<u>0.81</u>
<b>x</b>	<b>✓</b>	<b>✓</b>	6.954	0.049	<u>0.136</u>	0.186	0.82
<b>✓</b>	<b>✓</b>	<b>✓</b>	<b>6.833</b>	<b>0.046</b>	<b>0.130</b>	<b>0.176</b>	<b>0.79</b>

Table 3. **Ablation study on network and loss design.** The combination of  $\mathcal{L}_{\text{static}}$  and  $\mathcal{L}_{\text{cycle}}$  offers the best performance.

**Loss design.** We ablate the two additional loss terms  $\mathcal{L}_{\text{static}}$  and  $\mathcal{L}_{\text{cycle}}$  with the context network. The results are presented in Tab. 3, which demonstrate that the model trained

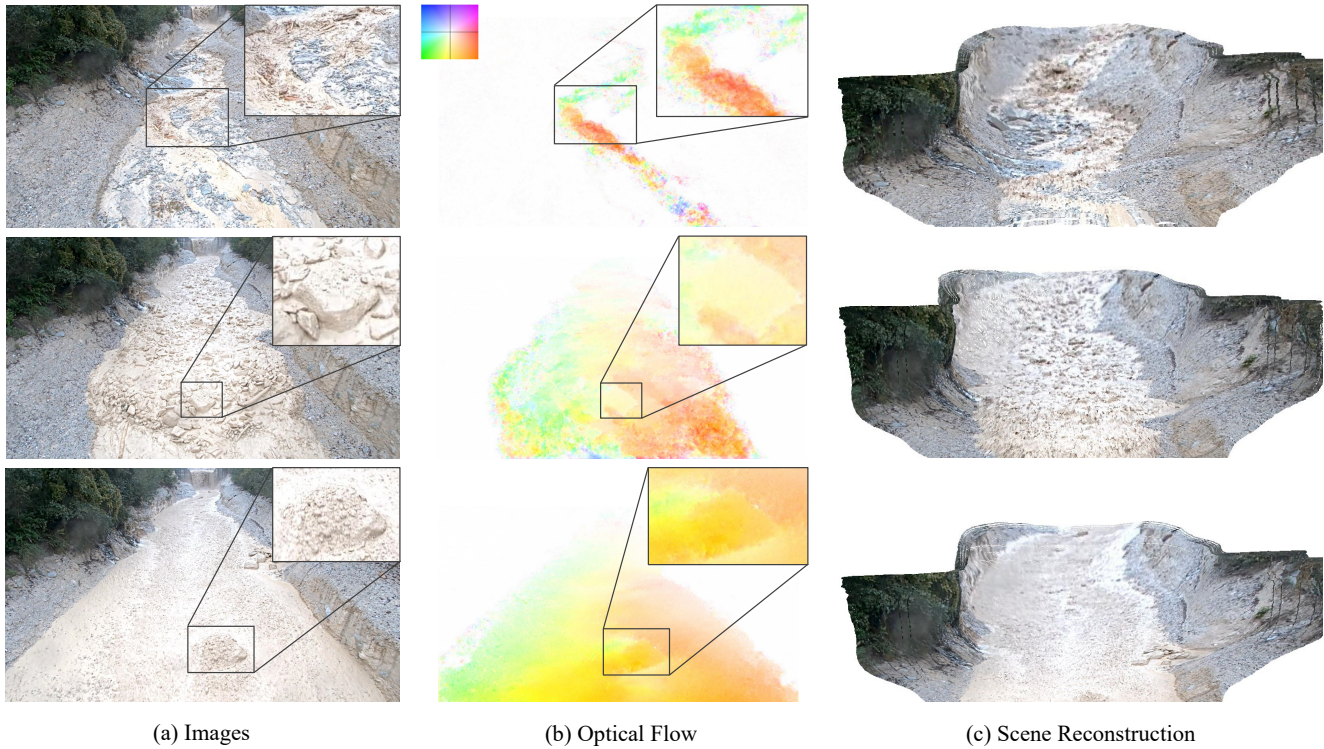


Figure 7. **Qualitative results.** Raw images (a), optical flow estimates (b), and scene reconstruction (c).

Ratio	Optical Flow↓	Depth Estimation↓			
	RMSD	MAE <sub>10</sub> [m]	MAE <sub>30</sub> [m]	MAE <sub>50</sub> [m]	Abs. Rel.[%]
0.2	<u>6.990</u>	0.057	0.196	0.269	0.11
0.5	7.018	<u>0.053</u>	<u>0.139</u>	<u>0.196</u>	<u>0.83</u>
0.8	<b>6.982</b>	<b>0.048</b>	<b>0.131</b>	<b>0.183</b>	<b>0.79</b>

Table 4. **Ablation study on downsampling ratios.**

with all loss terms performs the best, indicating that each of the losses contributes to overall model performance. We also observe that while the context network (Fig. 2) on its own is not particularly helpful, it significantly enhances the model performance when used in conjunction with our loss design.

**Depth downsampling.** To validate our downsampling strategy for depth supervision, we train three identical models using different depth input ratios and examine the resulting performance changes, as shown in Tab. 4. The results indicate that using a higher input ratio leads to a more precise depth map, suggesting that the network is able to leverage varying levels of depth information and adjust its performance accordingly.

## 5. Conclusion

In this study, we have focused on motion estimation under the unique properties exhibited by debris flows, and have presented a multi-task learning approach that integrates images and point clouds. Our self-supervised method is able to estimate the scene structure (depth) and the motion field at a level of detail that is unprecedented in the context of debris-flow research. Further analysis of these results will provide insights into the mechanisms that govern debris-flow motion, which can ultimately be used to reduce the risk associated with these destructive events. We release the dataset to the vision community, providing a new analysis domain and performance benchmark. For *future work*, one promising avenue is to incorporate physical constraints from fluid and soil mechanics to boost the model training. Another important direction is to develop automated algorithms for infrared cameras, in order to enable monitoring during nighttime.

**Acknowledgements.** Funding for this project was provided in part by an Swiss National Science Foundation grant to J. Aaron (p 193081). The authors further thank Stefan Boss, Christoph Graf, Brian McArdeall and Raffaele Spielmann for interesting discussions, as well as support in setting up and running the video camera and LiDAR systems.



## References

- [1] Jordan Aaron, Raffaele Spielmann, Brian W. McArdell, and Christoph Graf. Debris-flow monitoring with high-frequency LiDAR scanners: A new method to infer the internal dynamics of debris flows. In *Int'l Conf. on Debris Flow Hazard Mitigation*, 2023. 1, 2, 6, 7
- [2] Jordan Aaron, Raffaele Spielmann, Brian W. McArdell, and Christoph Graf. A novel method to investigate the dynamics of full-scale events in the field. *Geophysical Research Letters*, 2023. 1, 2, 6, 7
- [3] Simon Baker, Daniel Scharstein, JP Lewis, Stefan Roth, Michael J Black, and Richard Szeliski. A database and evaluation methodology for optical flow. *IJCV*, 92:1–31, 2011. 4, 6
- [4] Ramy Battrawy, René Schuster, Oliver Wasenmüller, Qing Rao, and Didier Stricker. Lidar-flow: Dense scene flow estimation from sparse lidar and stereo images. In *IROS*, 2019. 2
- [5] Aseem Behl, Despoina Paschalidou, Simon Donne, and Andreas Geiger. PointFlowNet: Learning representations for rigid motion estimation from point clouds. In *CVPR*, 2019. 1
- [6] Fabian Brickwedde, Steffen Abraham, and Rudolf Mester. Mono-SF: Multi-view geometry meets single-view depth for monocular scene flow estimation of dynamic traffic scenes. In *ICCV*, 2019. 2
- [7] D. J. Butler, J. Wulff, G. B. Stanley, and M. J. Black. A naturalistic open source movie for optical flow evaluation. In *ECCV*, 2012. 6
- [8] Holger Caesar, Varun Bankiti, Alex H Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nuScenes: A multi-modal dataset for autonomous driving. In *CVPR*, 2020. 2, 4
- [9] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected CRFs. *IEEE TPAMI*, 40(4):834–848, 2017. 4
- [10] Peter Constantin and Ciprian Foias. *Navier-Stokes equations*. University of Chicago Press, 2020. 4
- [11] Andreas Geiger, Philip Lenz, Christoph Stiller, and Raquel Urtasun. Vision meets robotics: The KITTI dataset. *IJRR*, 32(11):1231–1237, 2013. 2
- [12] Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for autonomous driving? the KITTI vision benchmark suite. In *CVPR*, 2012. 4
- [13] Clément Godard, Oisín Mac Aodha, and Gabriel J Brostow. Unsupervised monocular depth estimation with left-right consistency. In *CVPR*, 2017. 4
- [14] Clément Godard, Oisín Mac Aodha, Michael Firman, and Gabriel J Brostow. Digging into self-supervised monocular depth estimation. In *CVPR*, 2019. 4
- [15] Zan Gojcic, Or Litany, Andreas Wieser, Leonidas J Guibas, and Tolga Birdal. Weakly supervised learning of rigid 3d scene flow. In *CVPR*, 2021. 1, 2, 5
- [16] Christoph Holst, Jannik Janßen, Berit Schmitz, Martin Blome, Malte Dercks, Anna Schoch-Baumann, Jan Blöthe, Lothar Schrott, Heiner Kuhlmann, and Tomislav Medic. Increasing spatio-temporal resolution for monitoring alpine soilfluctuation using terrestrial laser scanners and 3d vector fields. *Remote Sensing*, 13(6):1192, 2021. 1
- [17] Shengyu Huang, Zan Gojcic, Jiahui Huang, Andreas Wieser, and Konrad Schindler. Dynamic 3d scene analysis by point cloud accumulation. In *ECCV*, 2022. 1, 2, 5
- [18] Oldrich Hungr, Serge Leroueil, and Luciano Picarelli. The Varnes classification of landslide types, an update. *Landslides*, 11:167–194, 2014. 1
- [19] Junhwa Hur and Stefan Roth. Self-supervised monocular scene flow estimation. In *CVPR*, 2020. 2, 3, 4, 6
- [20] Junhwa Hur and Stefan Roth. Self-supervised multi-frame monocular scene flow. In *CVPR*, 2021. 2, 4
- [21] Marcel Hürlimann, Velio Coviello, Coraline Bel, Xiaojun Guo, Matteo Berti, Christoph Graf, Johannes Hübl, Shusuke Miyata, Joel B Smith, and Hsiao-Yuan Yin. Debris-flow monitoring and warning: Review and examples. *Earth-Science Reviews*, 199:102981, 2019. 2
- [22] K Hutter, B Svendsen, and Rickenmann. Debris flow modeling: A review. *Continuum Mechanics and Thermodynamics*, 8(1):1–35, 1994. 1
- [23] Rico Jonschkowski, Austin Stone, Jonathan T Barron, Ariel Gordon, Kurt Konolige, and Anelia Angelova. What matters in unsupervised optical flow. In *ECCV*, 2020. 4
- [24] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 5
- [25] Jason Ku, Ali Harakeh, and Steven L Waslander. In defense of classical image processing: Fast depth completion on the CPU. In *CRV*, 2018. 6, 7
- [26] Katrin Lasinger, Christoph Vogel, Thomas Pock, and Konrad Schindler. 3d fluid flow estimation with integrated particle reconstruction. *IJCV*, 2020. 1
- [27] Katrin Lasinger, Christoph Vogel, and Konrad Schindler. Volumetric flow estimation for incompressible fluids using the stationary stokes equations. In *ICCV*, 2017. 1
- [28] Xueqian Li, Jhony Kaesemodel Pontes, and Simon Lucey. Neural scene flow prior. *NeurIPS*, 2021. 1
- [29] Haisong Liu, Tao Lu, Yihui Xu, Jia Liu, Wenjie Li, and Lijun Chen. CamLiFlow: Bidirectional camera-LiDAR fusion for joint optical flow and scene flow estimation. In *CVPR*, 2022. 2
- [30] Brian W McArdell and Mario Sartori. The Illgraben torrent system. *Landscapes and landforms of Switzerland*, pages 367–378, 2021. 6
- [31] Simon Meister, Junhwa Hur, and Stefan Roth. Unflow: Unsupervised learning of optical flow with a bidirectional census loss. In *AAAI*, 2018. 4, 5, 6
- [32] Hans-Hellmut Nagel and Wilfried Enkelmann. An investigation of smoothness constraints for the estimation of displacement vector fields from image sequences. *IEEE TPAMI*, (5):565–593, 1986. 6
- [33] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. PyTorch:

- An imperative style, high-performance deep learning library. *NeurIPS*, 2019. 5
- [34] David Petley. Global patterns of loss of life from landslides. *Geology*, 40(10):927–930, 2012. 1
- [35] René Ranftl, Katrin Lasinger, David Hafner, Konrad Schindler, and Vladlen Koltun. Towards robust monocular depth estimation: Mixing datasets for zero-shot cross-dataset transfer. *IEEE TPAMI*, 2020. 4
- [36] Rishav Rishav, Ramy Battrawy, René Schuster, Oliver Wasenmüller, and Didier Stricker. DeepLiDARFlow: A deep learning architecture for scene flow estimation using monocular camera and sparse lidar. In *IROS*, 2020. 2
- [37] Raffaele Spielmann, Jordan Aaron, , and Brian W. McARDell. Inferring spatial variations in velocity profiles and bed geometry of natural debris flows based on discharge estimates from high-frequency 3d LiDAR point clouds; Illgraben, Switzerland. In *Int'l Conf. on Debris Flow Hazard Mitigation*, 2023. 2, 6
- [38] Fridtjof Stein. Efficient computation of optical flow using the census transform. In *Joint Pattern Recognition Symposium*, pages 79–86, 2004. 6
- [39] Deqing Sun, Xiaodong Yang, Ming-Yu Liu, and Jan Kautz. PWC-net: Cnns for optical flow using pyramid, warping, and cost volume. In *CVPR*, 2018. 2, 3, 4
- [40] Tatsunori Tanai, Sudipta N Sinha, and Yoichi Sato. Fast multi-frame stereo scene flow with motion segmentation. In *CVPR*, 2017. 2
- [41] Zachary Teed and Jia Deng. RAFT: Recurrent all-pairs field transforms for optical flow. In *ECCV*, 2020. 6
- [42] Christoph Vogel, Stefan Roth, and Konrad Schindler. View-consistent 3d scene flow estimation over multiple frames. In *ECCV*, 2014. 2
- [43] Christoph Vogel, Konrad Schindler, and Stefan Roth. 3d scene flow estimation with a rigid motion prior. In *ICCV*, 2011. 2
- [44] Christoph Vogel, Konrad Schindler, and Stefan Roth. Piecewise rigid scene flow. In *ICCV*, 2013. 1
- [45] Chaoyang Wang, Xueqian Li, Jhony Kaesemodel Pontes, and Simon Lucey. Neural prior for trajectory estimation. In *CVPR*, 2022. 2
- [46] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE TIP*, 13(4):600–612, 2004. 4
- [47] George Wolberg. *Digital image warping*, volume 10662. IEEE Computer Society Press, 1990. 3
- [48] Alex Wong and Stefano Soatto. Unsupervised depth completion with calibrated backprojection layers. In *ICCV*, 2021. 4
- [49] Oliver Woodford, Philip Torr, Ian Reid, and Andrew Fitzgibbon. Global stereo reconstruction under second-order smoothness priors. *IEEE TPAMI*, 31(12):2115–2128, 2009. 4
- [50] Wenxuan Wu, Zhi Yuan Wang, Zhuwen Li, Wei Liu, and Li Fuxin. PointPWC-Net: Cost volume on point clouds for (self-) supervised scene flow estimation. In *ECCV*, 2020. 3
- [51] Haofei Xu, Jing Zhang, Jianfei Cai, Hamid Rezaatofghi, and Dacheng Tao. GMFlow: Learning optical flow via global matching. In *CVPR*, 2022. 5
- [52] Haofei Xu, Jing Zhang, Jianfei Cai, Hamid Rezaatofghi, Fisher Yu, Dacheng Tao, and Andreas Geiger. Unifying flow, stereo and depth estimation. *arXiv preprint arXiv:2211.05783*, 2022. 5
- [53] Gengshan Yang and Deva Ramanan. Upgrading optical flow to 3d scene flow through optical expansion. In *CVPR*, 2020. 2
- [54] Fisher Yu and Vladlen Koltun. Multi-scale context aggregation by dilated convolutions. In *ICLR*, 2016. 4
- [55] Mingrui Zhang, Jianhong Wang, James B Thomole, and Matthew Piggott. Learning to estimate and refine fluid motion with physical dynamics. In *ICML*, 2022. 1
- [56] Yuliang Zou, Zelun Luo, and Jia-Bin Huang. Df-net: Unsupervised joint learning of depth and flow using cross-task consistency. In *ECCV*, 2018. 6