

# SceneFlowFields++: Multi-frame Matching, Visibility Prediction, and Robust Interpolation for Scene Flow Estimation

René Schuster · Oliver Wasenmüller · Christian Unger · Georg  
Kuschik · Didier Stricker

Received: date / Accepted: date

**Abstract** State-of-the-art scene flow algorithms pursue the conflicting targets of accuracy, run time, and robustness. With the successful concept of pixel-wise matching and sparse-to-dense interpolation, we shift the operating point in this field of conflicts towards universality and speed. Avoiding strong assumptions on the domain or the problem yields a more robust algorithm. This algorithm is fast because we avoid explicit regularization during matching, which allows an efficient computation. Using image information from multiple time steps and explicit visibility prediction based on previous results, we achieve competitive performances on different data sets. Our contributions and results are evaluated in comparative experiments. Overall, we present an accurate scene flow algorithm that is faster and more generic than any individual benchmark leader.

**Keywords** Scene Flow · Matching · Occlusions · Interpolation

## 1 Introduction

Scene flow is the problem of estimating the perceived dense 3D motion field along with the 3D geometry of the scene. Typically, a set of stereo image pairs is used to estimate scene flow. Motion estimation is a crucial


---

R. Schuster · O. Wasenmüller · D. Stricker  
DFKI - German Research Center for Artificial Intelligence  
Kaiserslautern, Germany  
E-mail: firstname.lastname@dfki.de

C. Unger · G. Kuschik  
BMW Group  
Munich, Germany  
E-mail: christian.unger@bmw.de  
E-mail: g.kuschik@gmx.de



(a) Reference image with predicted invisibility (red).



(b) Final result of SceneFlowFields++.

Fig. 1: SceneFlowFields++ uses explicit visibility reasoning and the information of multiple frames to correctly estimate scene flow in occluded and out-of-bounds regions. Regularization is imposed by a consistency check and robust interpolation.

part for autonomous systems like self-driving vehicles and in many applications such as collision detection and path planning in driver assistance systems and robot navigation, object tracking, moving object detection,

frame interpolation (temporal super resolution), and many others.

Traditionally, motion is described by optical flow in 2D. Recently, with more computational resources and stereo cameras in mobile systems and vehicles available, the more complex, yet more realistic representation of motion in 3D is gaining increasing interest from research and industry. Despite growing efforts, either problem under unconstrained settings is far from being solved. However, with certain assumptions and restrictions to special domains, there exist methods that achieve impressive results on popular benchmarks. Yet, any assumption – even if it’s reasonable for a certain domain – will limit the applicability to other domains. The variance of performance of an algorithm on different data sets is one indicator for the robustness of that algorithm. Especially when applying deep neural networks, the generalization to different domains suffers heavily, unless the specific network is retrained or fine-tuned for the new domain. On the other hand, robustness alone should not be the only goal. A reasonable level of accuracy is required for any application that uses scene flow. Often, high accuracy comes at the price of run time. Top performing methods on the KITTI benchmark (Menze and Geiger, 2015) have run times of more than 5 minutes. Existing algorithms are focused on one or at most two of these targets, either speed, accuracy, or robustness. Upcoming deep learning approaches are potentially fast and accurate, but lack generalization. The state-of-the-art can be very accurate – sometimes even across different domains – but has high run times often.

With our novel concept of SceneFlowFields (SFF) (Schuster et al., 2018c), we want to push the limits of scene flow estimation towards a robust, comparatively fast, and accurate algorithm. Towards this end, we restrict ourselves to very few reasonable assumptions with minor impact on the versatility. Further, we avoid explicit regularization in our scene flow problem formulation which enables a fast computation of the solution. Regularization is imposed by outlier filtering and interpolation as it is often done in state-of-the-art optical flow algorithms (Bailer et al., 2015; Hu et al., 2017; Revaud et al., 2015; Schuster et al., 2018b). This approach is fast and domain invariant. However, the basis is accurate matching based on visual similarity in images. Therefore, we have to deal with the typical problems of matching tasks: Changes in lighting, image noise, perspective deformations, occlusions, etc. To further improve the accuracy and robustness of this sparse-to-dense pipeline, we extend the concept to a multi-frame setup with explicit visibility reasoning. We also use novel concepts for interpolation to increase the robustness.

Working in the context of classical, image-based scene flow estimation, we assume to have consecutive frame pairs of a stereo image sequence (at least 4 images, cf. Figure 2) as input. We represent scene flow as a 4D vector  $\mathbf{s} = (u, v, d_0, d_1)^T$  consisting of two optical flow components  $u, v$  and the disparity values  $d_0, d_1$  for both time steps. This way, we can express pixel correspondences in the four relevant images. Having rectified stereo image pairs and calibrated cameras, the 4D representation in image space is equivalent to the full scene flow representation of 3D geometry and 3D motion in world space.

This paper re-presents our work of SFF (Schuster et al., 2018c) and extends it considerably. In detail, the contributions of the original conference paper are:

- The first method of sparse-to-dense interpolation for scene flow.
- A novel method to find scene flow matches.
- A new interpolation method for scene flow that preserves boundaries of geometry and motion.
- An improved edge detector to approximate scene flow boundaries for the KITTI data set.
- An optional approach for straightforward integration of ego-motion.

We substantially extend SFF in this paper by the following contributions:

- Novel matching and filtering for scene flow in a multi-frame instead of dual-frame setup.
- Explicit one-shot visibility reasoning for scene flow matching on pixel-level.
- More robust two-stage interpolation of scene flow correspondences.
- A unified boundary detector for different data sets.
- Thorough evaluation of our conference approach and our novel extension with comparison to state-of-the-art on KITTI and MPI Sintel.

Since the extensions are based on the original SFF, we term our novel approach SceneFlowFields++ (SFF++).

The rest of the paper is structured as follows: We give an overview of related work in Section 2. Section 3 presents details about SFF, followed by an incremental description of SFF++ in Section 4. Experimental comparisons between SFF and SFF++ and to state-of-the-art are described in Section 5. We summarize our findings in Section 6.

## 2 Related Work

*Data Sets.* Since it is hard to capture ground truth scene flow information, there exist only very few data sets to evaluate scene flow algorithms on. Most of them use virtually rendered scenes to obtain the ground truth

data (Butler et al., 2012; Gaidon et al., 2016; Mayer et al., 2016). To the best of our knowledge the only realistic data set that provides a benchmark for scene flow is the KITTI Vision Benchmark (Geiger et al., 2012) that combines various tasks for automotive vision. Its introduction has played an important role in the development of stereo and optical flow algorithms, and the extension by Menze and Geiger (2015) has also driven the progress in scene flow estimation.

*Variational Scene Flow Estimation.* Vedula et al. (1999) were among the first to compute 3D scene flow. Afterwards, many variational approaches for scene flow estimation followed. First using pure color images as input (Basha et al., 2013; Huguet and Devernay, 2007) and later using RGB-D images (Herbst et al., 2013; Jaimez et al., 2015; Wedel et al., 2008). While a variational formulation is typically complex, Jaimez et al. (2015) achieved real-time performance with a primal-dual framework. Yet, all these approaches are sensitive to initialization and can not cope with large displacements, which is why they use a coarse-to-fine scheme. That in turn tends to miss finer details. Furthermore, the RGB-D approaches rely on depth sensors that either perform poorly in outdoor scenarios or are accordingly very expensive (Yoshida et al., 2017). Nowadays, variational methods are outperformed in terms of speed and accuracy by other approaches and are only used as a refinement step.

*Combination Approach.* Because the scene flow problem is highly related to the auxiliary tasks of optical flow and stereo disparity estimation, people have tried to estimate scene flow by combining separate results for optical flow and stereo disparity (Schuster et al., 2018a). Though the separation brings advantages for the complexity of the problem and thus the run time, it is believed that a single formulation of the problem yields more consistent scene flow results, which are more accurate and realistic.

*Assumption of Planarity and Rigidity.* Due to the advent of a piece-wise rigid plane model (Vogel et al., 2013), scene flow has recently achieved a boost in performance. The majority of top performing methods at the KITTI Vision Benchmark employ this model to enforce strong regularization (Behl et al., 2017; Lv et al., 2016; Menze and Geiger, 2015; Neoral and Šochman, 2017; Vogel et al., 2015). Vogel et al. (2014, 2015) encode this model by alternating assignment of each pixel to a plane segment and each segment to a rigid motion, based on a discrete set of planes and motions in view-consistent manner over multiple frames. Behl et al. (2017); Menze and Geiger (2015); Neoral and Šochman (2017) further

lower the complexity of the model by the assumption that a scene consists of very few independently moving rigid objects. Thus each plane segment only needs to be assigned to one object. All segments assigned to the same object share the same motion. By propagation of objects over multiple frames, Neoral and Šochman (2017) achieve temporal consistency for the model of Menze and Geiger (2015). Behl et al. (2017) use deep learning to obtain semantic object information a-priori. Lv et al. (2016) solve the pixel-to-plane assignment and the plane-to-motion assignment in a continuous domain. Despite the remarkable accuracy on KITTI, many of these methods are not applicable to domains with different characteristics. The rigid motion assumption is strongly violated by articulated gestures and other non-rigid motions that often occur in the Sintel data set (Butler et al., 2012). The assumption made by Behl et al. (2017); Menze and Geiger (2015); Neoral and Šochman (2017) that there are only a few independent dynamic objects in a scene is inappropriate for highly dynamic scenarios. Further, methods falling into this category typically have very long run times of several minutes up to almost one hour per frame.

*Guided by Semantics.* Other scene flow algorithms use deep learning to incorporate semantic information into the motion estimation problem (Ren et al., 2017). Yet, in terms of robustness, deep learning approaches (Behl et al., 2017; Ren et al., 2017) typically lag behind because they generalize badly to unseen data and even worse to data from different domains (Wang and Deng, 2018). This is especially true for semantic segmentation, where the domain gap is amplified by the mismatch of semantic classes between domains (Lv et al., 2018). Our algorithm is especially designed to achieve good results across different data sets without any dedicated parameter tuning.

*Static-Dynamic Decomposition.* Yet another promising strategy builds on the decomposition of a scene into static and moving parts (Taniai et al., 2017). While the motion of dynamic objects is estimated by solving a discrete labeling problem (as by Chen and Koltun (2016)) using the Semi-Global-Matching (SGM) (Hirschmüller, 2008) algorithm, the perceived motion of all static parts is directly obtained from the 3D geometry of the scene and the ego-motion of the camera. This approach is especially convenient for scenes, where only a small proportion consists of moving objects, like it is usually the case in traffic scenarios. However, any a-priori assumption limits the versatility of a method. A rigid plane model performs poorly when applied to deformable objects, and ego-motion estimation for highly dynamic scenes is hard.

Table 1: Comparison of different categories of scene flow algorithms.

Category	Advantages	Disadvantages	Examples
Variational		Slow, inaccurate, only indoors	(Huguet and Devernay, 2007) (Herbst et al., 2013)
Decomposition	Potentially fast	Ego-motion dependency, inconsistent	(Schuster et al., 2018a) FSF (Tani ai et al., 2017)
Pice-wise Rigid Planes Model	Strong regularization	Slow	PRSF (Vogel et al., 2013) OSF (Menze and Geiger, 2015) CSF (Lv et al., 2016)
Deep Learning	Fast	Poor generalization	PWOC-3D (Saxena et al., 2019) DRISF (Ma et al., 2019)
Sparse-to-Dense	Comparatively fast, good generalization	Sensitive to distribution of matches	SFF Schuster et al. (2018c) SFF++ (ours)

*From Dual to Multiple Frames.* When speaking about the related work, one has to differentiate between dual-frame (Lv et al., 2016; Menze and Geiger, 2015; Vogel et al., 2013) and multi-frame (Neoral and Šochman, 2017; Tani ai et al., 2017; Vogel et al., 2015) approaches. Especially in the context of traffic scenarios, we have several characteristics that make matching between two frame pairs much more challenging than in a multi-view setting. These characteristics are: 1.) Considerably large stereo and flow displacements. 2.) Difficult lighting conditions and many reflective and (semi-)transparent surfaces of cars. 3.) Fast ego-motions sometimes combined with low frame rates, which causes large regions to move out of image bounds. This explains why there exist pairs of dual-frame and multi-frame methods. The transitions from OSF (Menze and Geiger, 2015) to OSF+TC (Neoral and Šochman, 2017) and from PRSF (Vogel et al., 2013) to Piece-wise Rigid Scene Model (PRSM) (Vogel et al., 2014, 2015) brought essential improvements by using the additional information of multiple image frames. However, this additional information comes at a cost. The relationship between multiple temporal steps needs to be modeled to make use of the additional images. A typical model is to assume smooth, constant motion between neighboring time steps (Neoral and Šochman, 2017; Vogel et al., 2015). In this extension of SFF (Schuster et al., 2018c), we also follow this trend and use more than two image pairs, assuming constant motion.

*Discrimination from State-of-the-Art.* We can cluster related work into certain categories (cf. Table 1). Because our scene flow method follows the newly introduced sparse-to-dense approach, it differs from any of the related approaches. We find sparse scene flow matches that are interpolated to a dense scene flow field, recovering the geometry of the scene and the 3D motion. Our

method has to be distinguished from purely variational approaches. Although we use variational optimization, it can be considered as a post-processing step for refinement. Similar, the cameras ego-motion is only used to refine the results for static image regions. During interpolation, we assume that the geometry of a scene can be modeled by very small planar segments, but we do not initially presume any coarse segmentation. In fact, the very small size of our plane segments leads to smoothly curved shapes and sharp boundaries. The same holds for our piece-wise motion model that is used to interpolate the 3D motion. Methods which are guided by semantic segmentation from deep neural networks will generalize badly to other domains, unless they are fine-tuned for the new task. Same is assumed for upcoming purely learning based approaches (Ma et al., 2019; Saxena et al., 2019) which are potentially even faster than our approach. SFF++ focuses especially on robustness across domains and applications.

We contrast the different categories and compare their properties in Table 1. The piece-wise rigid planes models is particularly accurate due to its strong regularization (as long as the assumptions are not violated), but is also complex and computationally expensive. The decomposition (separation) approach that splits the scene flow problem into less difficult sub-problems is especially fast and benefits from advances in the auxiliary tasks. Yet, separate computation leads to overall less consistent scene flow. Deep learning is potentially fast due to the inherent parallelization on GPUs, but sensitive to the distribution of available training data and not interpretable in case of failure. The novel scene flow concept of sparse-to-dense interpolation allows to separate matching from regularization. With the use of appropriate interpolation models and interpolation regions, the negative impact of violated assumptions

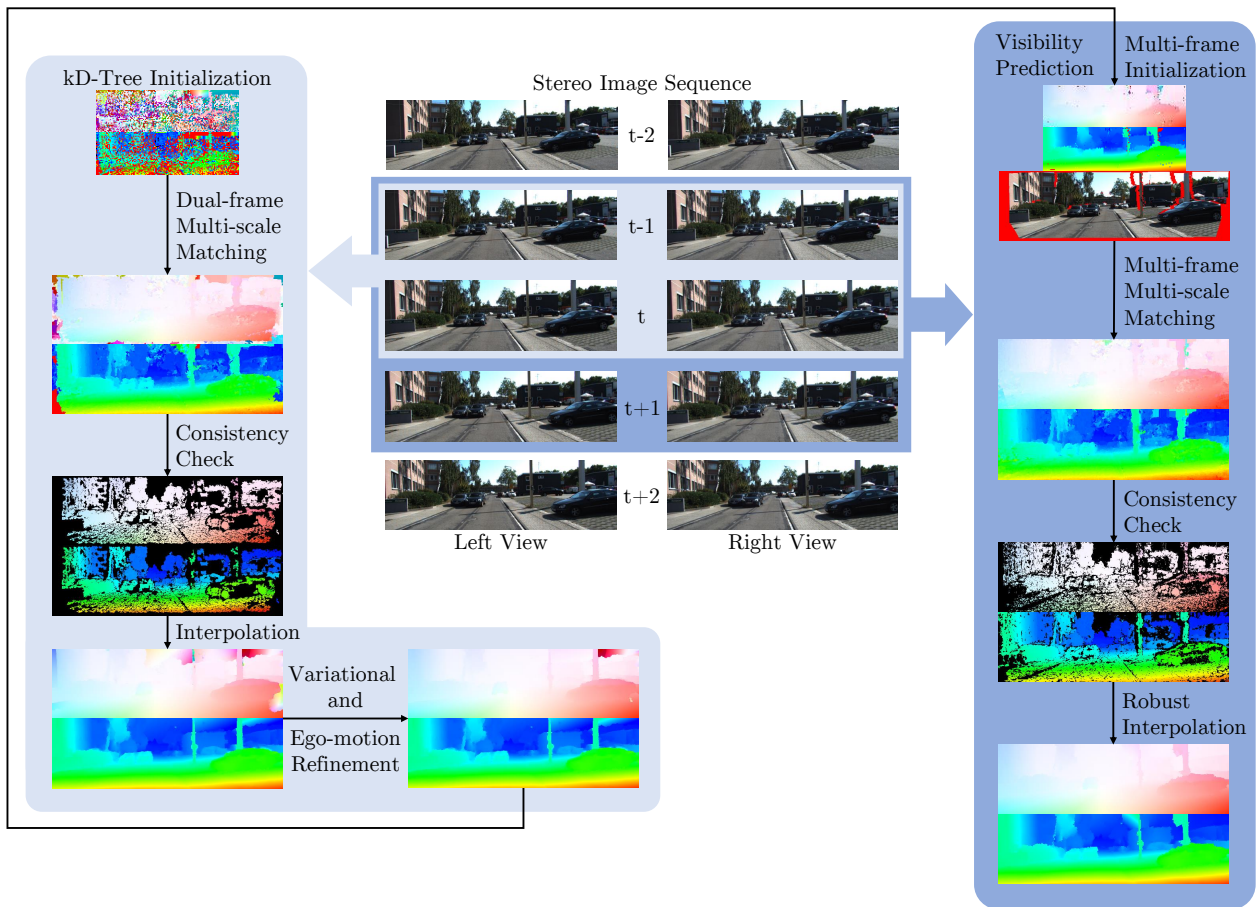


Fig. 2: Overview of our dual- and multi-frame approach. We work on a sequence of stereo image pairs. SFF (left) uses two frame pairs, initializes on a sub-scale and applies multi-scale propagation and random search for matching. The matched result is filtered in a consistency check, interpolated and refined by variational optimization and the optional ego-motion model. The final result is used to initialize the multi-frame process SFF++ (right) which uses three frame pairs. Matching is done with explicit visibility reasoning over all images. The consistency check is adjusted to the novel setup, and the interpolation uses novel concepts for increased robustness. Note that there is no refinement necessary in our improved robust multi-frame pipeline. Each step is illustrated by the corresponding optical flow and one disparity map of the respective scene flow field. Table 2 summarizes the differences between the dual-frame and multi-frame approach.

can be diminished. However, the separation of matching and regularization makes our approach sensitive to the quality of the sparse matching results. To overcome this issue, we propose to use multiple frames within this work.

### 3 SceneFlowFields

For our dual-frame scene flow estimation we assume to have the typical stereo image information provided, i.e. two rectified stereo image pairs  $(I_0^l, I_1^l, I_0^r, I_1^r)$  at times  $t$  and  $t + 1$  along with the camera intrinsics. We further assume that the baseline  $B$  is known. For rectified images, the baseline describes the relative pose between

the left and right cameras as translation parallel to the image plane. During matching, we jointly optimize all four components to obtain coherent scene flow. Our matching algorithm follows the ideas of non-regularized, coarse-to-fine matching with propagation and random search (Bailer et al., 2015, 2018) which was shown to work the best for sparse-to-dense optical flow pipelines (Bailer et al., 2015; Schuster et al., 2018b; Zweig and Wolf, 2017). Given the mentioned information, we estimate a dense scene flow field as follows: For  $k$  sub-scales we initialize the coarsest scale. For all  $k + 1$  scales (the  $k = 3$  sub-scales plus full resolution), we iteratively propagate scene flow vectors and adjust them by random search. Afterwards, the dense scene flow map on

full resolution is filtered using a two-stage consistency check. The filtered scene flow map is further thinned out in a sparsification step. Scene flow boundaries are detected using a structured random forest. Geometry and 3D motion are separately interpolated with respect to object boundaries of the scene. Finally, we refine the 3D motion by variational optimization. An overview of the dual-frame method SFM is outlined in [Figure 2](#).

### 3.1 Sparse Correspondences

*Matching Cost.* The matching cost in our algorithm solely depends on a data term. No additional smoothness assumptions are made like e.g. by [Herbst et al. \(2013\)](#); [Huguet and Devernay \(2007\)](#); [Lv et al. \(2016\)](#); [Menze and Geiger \(2015\)](#); [Vogel et al. \(2013, 2015\)](#). Given a scene flow vector, we define its matching cost by the sum of Euclidean distances between SIFTFlow features ([Liu et al., 2011](#)) over small image patches. Our matching error for two corresponding pixels  $\mathbf{p}$  and  $\mathbf{p}'$  in images  $I$  and  $I'$  is defined by the following cost

$$C(I, \mathbf{p}, I', \mathbf{p}') = \sum_{\tilde{\mathbf{p}} \in W(\mathbf{p})} \|\phi(I, \tilde{\mathbf{p}}) - \phi(I', \mathbf{p}' + \tilde{\mathbf{p}} - \mathbf{p})\|, \quad (1)$$

where  $W(\mathbf{p})$  is a  $7 \times 7$  patch window centered at pixel  $\mathbf{p}$  and  $\phi(I, \mathbf{p})$  is a function that returns the first three principal components of a SIFT feature vector (SIFT-Flow) for pixel  $\mathbf{p}$  in image  $I$ . We evaluate the cost for three image correspondences. The stereo image pair at time  $t$ , the temporal image pair for the left view point (standard optical flow correspondence) and a cross correspondence between the reference frame and the right frame at the next time step (cf. [Figure 2](#)). This leads to the following overall cost  $C$  for a scene flow vector  $\mathbf{s} = (u, v, d_0, d_1)^T$  at pixel  $\mathbf{p}$ :

$$C(\mathbf{p}, \mathbf{s}) = C\left(I_0^l, \mathbf{p}, I_1^l, \mathbf{p} + (u, v)^T\right) + C\left(I_0^l, \mathbf{p}, I_0^r, \mathbf{p} + (-d_0, 0)^T\right) + C\left(I_0^l, \mathbf{p}, I_1^r, \mathbf{p} + (u - d_1, v)^T\right). \quad (2)$$

With that we can obtain a dense scene flow field  $\mathcal{S}$  by optimizing the following energy minimization problem

$$E(\mathbf{I}, \mathcal{S}) = \int_{\Omega} C(\mathbf{p}, \mathcal{S}(\mathbf{p})) d\mathbf{p}. \quad (3)$$

$$\hat{\mathcal{S}} = \arg \min_{\mathcal{S}} E(\mathbf{I}, \mathcal{S}). \quad (4)$$

Though this optimization includes a lot of variables, we can exploit the fact that we can optimize

$\mathbf{s} = (u, v, d_0, d_1)^T$  for each pixel individually since our formulation includes no inter-pixel dependencies, e.g. no explicit regularization. Therefore, we use an efficient, greedy, stochastic optimization approach of propagation and random search.

*Initialization.* Initialization is based on kD-trees similar to the work of [He and Sun \(2012\)](#), but with three trees, using Walsh-Hadamard-Transform (WHT) features as done by [Bailer et al. \(2015\)](#); [Schuster et al. \(2018b\)](#); [Wannenwetsch et al. \(2017\)](#). For each frame other than the reference frame ( $I_0^l$ ), we compute a WHT feature vector per pixel and store them in a kD-tree. To initialize a pixel of the reference image, we query the feature vector of the pixel to the pre-computed kD-trees. Scene flow matches are then obtained by comparing all combinations of the leafs for each queried node according to the matching data term introduced before ([Equation 2](#)). Since our stereo image pairs are rectified, for the images observed from the right camera view, we create kD-trees which regard the epipolar constraint, i.e. queries for such a tree will only return elements, which lie on the same image row as the query pixel. This way, we can efficiently lower the number of leaves per node for the epipolar trees, which speeds up the initialization process without loss of accuracy. For further acceleration, we use this initialization on the coarsest resolution only, and let the propagation fill the gaps when evolving to the next higher scale.

Though the result of the initialization looks very noisy and unreliable (cf. [Figure 2](#)), it is much better than random initialization. In fact, we could not reproduce similar results of our full approach with random initialization or any other more simple initialization method. This is because although the WHT features are not very expressive, they are sufficient to find at least a few close-to-correct initial scene flow values, which later are spread and refined.

*Multi-Scale Propagation.* The initial matches are spread by propagation and steadily refined by random search. This is done over multiple scales which helps to distribute rare correct initial matches over the whole image. For each scale, we run several iterations of propagation in one out of the four image quadrants, so that each direction is used equally often. During propagation, a scene flow vector will be replaced if the propagated vector has a smaller matching cost (see [Equation 2](#)). If this is not the case, the propagation along this path will continue with the previously existing scene flow vector. After each iteration we perform a random search. This means that for all pixels we add a uniformly distributed random offset in the interval  $]-1, 1[$  in pixel

units of the current scale to each of the four scene flow components and check whether the matching cost decreases. Both propagation and random search help to obtain a smoothly varying vector field and to find correct matches even if the initialization is flawed. For the different scale spaces, we simulate the scaling by smoothing the images and taking only every  $n$ -th pixel for a subsampling factor of  $n = 2^k$ , so that the patches consist of the same number of pixels for all scales. This way, we prevent (up)sampling errors because all operations are performed on exact pixel locations on the full image resolution. Smoothing is done by area-based downsampling followed by upsampling using Lanczos interpolation. Note that a similar matching strategy has already been used by Bailer et al. (2015), but while they use it for optical flow, we apply it to twice as many dimensions in search space. This matching method was shown to work very well for motion fields. Even unique initial values are spread across the whole image yielding smooth motion fields with sharp boundaries.

*Consistency Check.* The matching procedure yields a dense map of scene flow correspondences across all images. However, many of the correspondences are wrong because of occlusions, out-of-bounds motion or simply because of mismatching due to challenging image conditions. To remove these outliers, we perform a two-step consistency check. First, we compute an inverse scene flow field for which the reference image is the right image at time  $t + 1$ . Temporal order as well as points of view are swapped. Everything else remains as explained above.

During consistency check, optical flow and both disparity maps for each pixel are compared to the corresponding values of the inverse scene flow field. If either difference exceeds a consistency threshold  $\tau_c = 1$  in image space (pixels), the complete scene flow vector will be removed. In a second stage, we form small regions of the remaining pixels as proposed by Bailer et al. (2015), where a pixel is added to a region if it has approximately the same scene flow vector. Afterwards, we check whether it is possible to add one of the already removed outliers in the neighborhood following the same rule. If this is possible and the region is smaller than  $s_c = 100$  pixels, we will remove the whole region. This way we obtain the filtered final scene flow correspondences of high accuracy and very few outliers (cf. Table 4).

Since we use a single scene flow formulation during consistency check also, the filtered results are very coherent. However, the joint filtering of all four scene flow components removes entire matches which might be only partly inconsistent. This applies especially to the reference disparity  $d_0$ . Since our interpolation model




Fig. 3: Sparse correspondences (left) and dense interpolation (right). Optical flow (a) and disparities at  $t$  (b) and  $t + 1$  (c).

(see Section 3.2) is split into separate steps for geometry and motion, we counter this problem, and fill up some gaps of the filtered disparity with additional values. These values are the result of a second independent consistency check for the disparity matches only. For the separate check we compute a second disparity map with SGM (Hirschmüller, 2008) and use the same threshold  $\tau_c$  as before. The additional disparity values that are retrieved this way are as accurate as the one from our standard consistency check, but much denser, which is shown in Figure 3. We merge the result of the second consistency check for disparity only with the result of the first full consistency check.

### 3.2 Dense Interpolation

*Sparsification.* Before interpolating the filtered scene flow field to recover full density, an additional sparsification step is performed. This helps to extend the spatial support of the neighborhood during interpolation and speeds up the whole process (Bailer et al., 2015). It also serves as a final filtering step, because for each non-overlapping  $3 \times 3$  block, we select the match with the lowest consistency error during filtering only. The remaining matches are called seeds with respect to the interpolation.

*Interpolation Boundaries.* A crucial part of the interpolation is the estimation of scene flow boundaries. While (Bailer et al., 2015; Revaud et al., 2015) approximate motion boundaries for optical flow with a texture-agnostic edge detector (Dollár and Zitnick, 2013), our edge detector is trained on semantic boundaries. We find that this models geometric boundaries, as well as motion boundaries, much better than image edges and is much more robust to lighting, shadows, and coarse textures. To do so, we have gathered 424 images of the KITTI data set from (Ošep et al., 2016; Ros et al., 2015; Xu et al., 2016)

that have been densely labeled with semantic class information. Within these images we have merged semantic classes that in general neither align with geometric nor motion discontinuities, e.g. lane markings and road, or pole and panel. The boundaries between the remaining semantic labels are used as binary edge maps to train our edge detector. To this end, we utilize the framework of Structured Edge Detection (SED) (Dollár and Zitnick, 2013) and train a random forest with the same parameters as in their paper, except for the number of training patches. We sample twice as many positive and negative patches during training because we use a bigger data set with images of higher resolution. The impact of the novel boundary detector will be evaluated in our ablation study in Section 5.1 and is visualized in Figure 4.

*Interpolation Models.* We interpolate geometry and 3D motion separately. This allows us to incorporate domain knowledge into the interpolation process by having dedicated models for both types of interpolation. Also, due to the separate consistency check for disparity only and full scene flow, we have different amounts of matches for geometry and motion, which favors the separation. Apart from the models and number of seeds, the general concept of interpolation is the same. Suppose a local, boundary-aware neighborhood of 160 and 80 seeds is given for each unknown scene flow vector  $\hat{\mathbf{s}}$  at pixel  $\hat{p}$  for geometric and motion seeds respectively,  $\mathcal{N}_{geo}$  and  $\mathcal{N}_{motion}$ . The depth of pixel  $\hat{p}$  is reconstructed by fitting a plane  $E(\hat{p}) : a_1x + a_2y + a_3 = d_0$  through all seeds of the neighborhood  $\mathcal{N}_{geo}$ . This is done by solving a linear system of equations for all neighboring seed points  $p_g$  for which the disparity values are known, using weighted least squares. The weights for each seed are obtained from a Gaussian kernel  $g(D) = \exp(-\alpha D)$  on the distance  $D(\hat{p}, p_g)$  between target pixel and seed with coefficient  $\alpha = 2.2$ . The missing disparity value of  $\hat{p}$  is obtained by plugging the coordinates of  $\hat{p}$  into the estimated plane equation. In a similar fashion, but using a neighborhood of motion seeds  $\mathcal{N}_{motion}$ , the missing 3D motion is obtained by fitting an affine 3D transformation  $\mathbf{x}_1 = \mathbf{A}\mathbf{x}_0 + \mathbf{t}$  using weighted least squares on all neighboring motion seeds  $p_m$ . Where  $\mathbf{x}_t = (x_t, y_t, z_t)^T$  are the 3D world coordinates of motion seed  $p_m$  at time  $t$  and  $t+1$ , and  $[\mathbf{A}|\mathbf{t}] \in \mathbb{R}^{3 \times 4}$  is the affine 3D transformation of twelve unknowns. The weights are computed by the same Gaussian kernel as for geometric interpolation, but using the distances  $D(\hat{p}, p_m)$  between the target pixel and the motion seeds. To summarize, for the full reconstruction of scene flow  $\hat{\mathbf{s}} = (u, v, d_0, d_1)^T$  at pixel  $\hat{p}$ , we compute  $d_0$  using the plane model  $E(\hat{p})$ , reproject the point into 3D world space, transform it according to

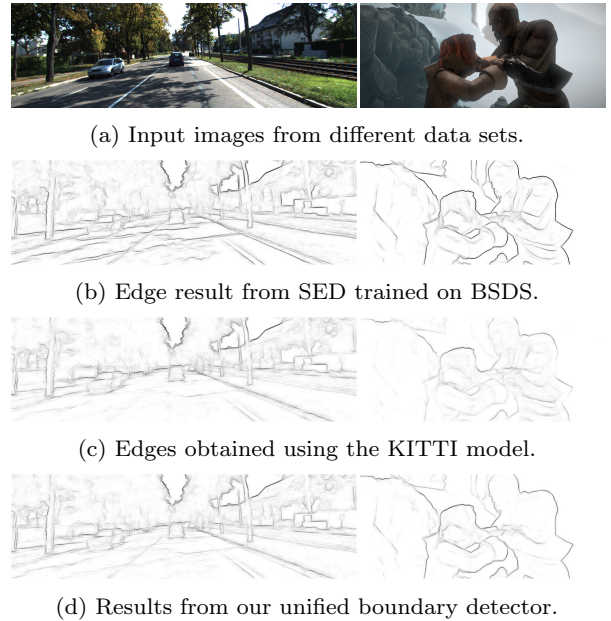


Fig. 4: Whereas SED (Dollár and Zitnick, 2013) (b) detects all image gradients, our KITTI boundary detector (c) suppresses lane markings and shadows. Our unified detector of Section 4.2 (d), achieves a good trade-off between the advantages of (c) and generalization abilities across different domains.

its associated affine transformation  $[\mathbf{A}|\mathbf{t}]$ , and project it back to image space to obtain  $u, v$  and  $d_1$ .

*Edge-Aware Neighborhood.* To find the local neighborhoods, we follow the idea of Revaud et al. (2015) using both of their approximations. That is first, the  $n$  closest seeds to a pixel  $\hat{p}$  are the  $n-1$  closest seeds to the closest seed of  $\hat{p}$ , thus all pixels with the same closest seed share the same local neighborhood. And secondly, the distance between  $\hat{p}$  and its closest seed is a constant offset for all neighboring seeds, which can be neglected. It is therefore sufficient to find a labeling that assigns each pixel to its closest seed and to find the local neighborhood for each seed. We use the graph-based method of Revaud et al. (2015) for this. The graph is constructed using the geodesic distances between seeds, that are directly based on the edge maps from our boundary detector. A strong boundary, i.e. a high value in the edge map, indicates a high cost for crossing that pixel, which leads to a high geodesic distance. The algorithm finds the shortest path between seeds and sums up the edge values along this path to obtain the geodesic distances between seeds.



### 3.3 Variational Optimization

To further refine the 3D motion after interpolation, we use variational energy minimization to optimize the objective

$$E(u, v, d') = E_{data}^{flow} + E_{data}^{cross} + \varphi \cdot E_{smooth}. \quad (5)$$

Motion is represented in image space by optical flow and the change in disparity  $d'$ . The energy consists of three parts. Two data terms, one temporal correspondence and one cross correspondence, and an adaptively weighted smoothness term for regularization. The data terms use the gradient constancy assumption. Our experiments have shown, that a term for the color constancy assumption can be neglected.

$$E_{data}^*(I, I', \mathbf{x}, \mathbf{w}) = \int_{\Omega} \beta(\mathbf{x}, \mathbf{w}) \cdot \Psi\left(\gamma \cdot |\nabla I'(\mathbf{x} + \mathbf{w}) - \nabla I(\mathbf{x})|^2\right) dx \quad (6)$$

with  $\mathbf{w}$  being a placeholder for either the optical flow displacement in Equation 7 or the cross image displacement in Equation 8.

$$E_{data}^{flow} = E_{data}^*(I_0^l, I_1^l, \mathbf{x}, (u, v)^T) \quad (7)$$

$$E_{data}^{cross} = E_{data}^*(I_0^l, I_1^r, \mathbf{x}, (u - d_0 - d', v)^T) \quad (8)$$

The data terms do not contribute to the energy if the function

$$\beta(\mathbf{x}, \mathbf{w}) = \begin{cases} 1, & \text{if } (\mathbf{x} + \mathbf{w})^T \in \Omega \\ 0, & \text{otherwise} \end{cases} \quad (9)$$

indicates that the scene flow is leaving the image domain (out-of-bounds motion). The smoothness term

$$E_{smooth} = \int_{\Omega} \Psi\left(|\nabla u|^2 + |\nabla v|^2 + \lambda \cdot |\nabla d'|^2\right) dx \quad (10)$$

penalizes changes in the motion field where  $|\nabla u|^2$  is the magnitude of the spatial derivatives of one scene flow component. It is weighted by

$$\varphi(\mathbf{x}) = e^{-\kappa B(\mathbf{x})}, \quad (11)$$

where  $B(\mathbf{x})$  is the edge value of our boundary detector at pixel  $\mathbf{x}$  (cf. Section 3.2) and  $\kappa = 5$  is a kernel coefficient. All parts use the Charbonnier penalty  $\Psi(x^2) = \sqrt{x^2 + \varepsilon^2}$  with  $\varepsilon^2 = 0.0001$  to achieve robustness. Since the smoothness term rather enforces constancy instead of smoothness if  $\beta$  for both data terms is zero, we do not optimize the scene flow at pixels where the interpolated scene flow field leaves  $\Omega$ . Our energy formulation is inspired by (Brox et al., 2004; Huguet




Fig. 5: Example of our motion segmentation. Sparse motion indicators as obtained during ego-motion computation (a), dense segmentation by interpolation (b) and moving ground truth objects as provided by KITTI (Menze and Geiger, 2015) (c).

and Devernay, 2007; Wedel et al., 2008). We use linear approximations of the Euler-Lagrange equations for the objective in Equation 5 and apply the framework of Brox et al. (2004) without the coarse-to-fine steps to find a solution by Successive Over-Relaxation (SOR). Again, we will compare the results with and without the variational refinement in our ablation study in Section 5.1.

### 3.4 Optional Ego-Motion Model

In Section 5 we will show that our approach as described so far achieves competitive results. For the special challenges of the KITTI data set, we make an additional, optional assumption to further improve the performance of SFF. Following Taniai et al. (2017), we argue that most parts of a scene are static and thus that the 3D motion for these areas is fully determined by the ego-motion of the observer. Given the ego-motion and a motion segmentation into static and dynamic areas, we apply the inverse ego-motion to all static points in the scene. Using our matching and interpolation scheme, both can be estimated easily with almost no additional effort. However, our improvements in Section 4 will make the refinement step by the ego-motion model obsolete. Our experiments will show, that the impact of the ego-motion optimization diminishes when using our more robust interpolation.

*Ego-Motion Estimation.* The filtered scene flow field before interpolation provides very accurate matches across

all images. We compute 3D-2D correspondences between the reference frame and the temporally subsequent frame by triangulation with the stereo matches. We limit the depth of these correspondences to 35 meters because disparity resolution for farther distances becomes too inaccurate. This way, we obtain a Perspective-n-Point (PnP) problem, which we solve iteratively using Levenberg-Marquardt and RANSAC to find the relative pose between the left camera at time  $t$  and  $t + 1$  by minimizing the re-projection error of all correspondences. For RANSAC, we consider a correspondence to be an outlier if the re-projection error is above 1 pixel. After a first estimation, we recompute the set of inliers with a relaxed threshold of 3 pixels and re-estimate the pose  $P = [R|t] \in \mathbb{R}^{3 \times 4}$ . The two-stage process helps to avoid local optima and to find a trade-off between diverse and robust correspondences.

*Motion Segmentation.* An initial sparse motion segmentation can directly be obtained as a side product of the ego-motion estimation. Outliers in RANSAC are considered *in motion*, while points in conformity with the estimated ego-motion are marked as *static*. We use our boundary-aware interpolation to compute a dense segmentation (cf. Figure 5). Pixels labeled as moving are spread up to the object boundaries within which they are detected. Because the segmentation is only a binary labeling, no complex interpolation model is needed. An unknown pixel gets assigned with the weighted mean of its local neighborhood. The weights are again based on the geodesic distances between matches. This interpolation method is similar to the Nadaraya-Watson estimator described by Revaud et al. (2015). The interpolated motion field is then thresholded to obtain a dense, binary motion segmentation. Finally, the inverse estimated ego-motion is applied to all points that are labeled as static. The impact of the ego-motion model is huge for our dual-frame conference approach, as will be shown in Section 5.1.

#### 4 Robust Multi-frame Extension

Our experiments will show that SFF, presented in the previous section, yields competitive scene flow results, especially when the ego-motion model is applied. However, there are two main problems with the presented pipeline: 1.) Pixel-wise matching without regularization is error prone under some circumstances (saturation, lighting variations, homogeneous or repetitive textures, etc.) and even impossible in occluded and other invisible image regions (e.g. out-of-bounds motions, cf. Figure 7a). 2.) The accuracy of the interpolation suffers from increasing

Table 2: A comparison between the dual- and multi-frame pipeline. We contrast the individual steps of the pipeline. The differences are visualized in Figure 2.

	<b>SFF</b> dual-frame	<b>SFF++</b> multi-frame
Initialization	kD-Trees	Previous result
Matching	Two time steps	Three time steps and visibility reasoning
Consistency Check	Fully inverse	Left-right
Interpolation	Edge-preserving	Edge-preserving and robust
Refinement	Variational and ego-motion	Not needed

gap sizes in the filtered scene flow field and from remaining outliers after consistency check. The ego-motion model can compensate inaccurate interpolation to some extent, but not sufficiently. In this paper, we address both major issues with our novel robust extension to multiple frame pairs.

There are different techniques to handle the unmatchable parts of a scene, i.e. cases where it is difficult or even impossible to find correspondences in the relevant images. Typically, some kind of regularization is applied, like in form of a smoothness assumption which encodes that neighboring pixels should represent a similar motion so that local visual evidence can support the motion estimation in the difficult areas. In methods that employ the piece-wise rigid plane model (Behl et al., 2017; Lv et al., 2016; Menze and Geiger, 2015; Neoral and Šochman, 2017; Vogel et al., 2013, 2015) this kind of regularization is for two reasons considerably strong. Firstly, each local patch describing a slanted plane undergoes the same transformation by design. Secondly, inter-plane smoothness is further enforced by dedicated terms in the energy formulation. However, regularization terms increase the computational effort significantly (cf. Table 5), and would prohibit the use of our efficient optimization strategy of Equation 3. An alternate concept to handle unmatchable regions is sparse-to-dense interpolation. This idea is rather young and was first successfully realized by EPICFlow (Revaud et al., 2015) for the optical flow problem. The idea is to remove regions of low confidence (i.e. regions where regularization would be required to match them accurately) and to use interpolation to fill the gaps based on reliable matches. SFF is to the best of our knowledge the first method of sparse-to-dense interpolation for scene flow correspondences and our extension is transferring this concept to a multi-frame setup for the first time. Even though the consistency check removes outliers reliably, the gaps can not be refilled by the interpolation correctly in some scenarios. Therefore, we tackle the problem before it occurs by using image information from multiple frames to

avoid mismatches and resolve ambiguity in unmatchable regions. This will result in more accurate and better distributed matches of higher density as will be shown in Section 5.2. Further, we transfer the concepts for interpolation from RICFlow (Hu et al., 2016) to the scene flow domain to improve robustness during interpolation.

All our extensions using more than two stereo image pairs build on the assumption of constant motion. That means, we assume the observed motion at both time steps (from  $t-1$  to  $t$  and from  $t$  to  $t+1$ ) to be the same. The error made by this assumption converges towards zero for continuous motions with increasing frame rates. Our multi-frame approach is additionally designed to process video streams in an online manner. That means that the first two frame pairs of a sequence are processed with the dual-frame approach from Section 3. All subsequent frames are then added in an incremental way and processed within a sliding temporal window of three frame pairs. We exploit the additional information of the extra images during matching in three ways. First, the previous results are used during initialization. Secondly, we use the previous scene flow to predict visibility of the scene. Lastly, we use our assumption together with the visibility prediction to match scene flow across all six relevant images.

The overall structure of our multi-frame approach remains the same as in the dual-frame method but with two more images. Using a set of six input images  $\mathbf{I} = \{I_0^l, I_0^r, I_1^l, I_1^r, I_{-1}^l, I_{-1}^r\}$ , we find accurate matches, remove possible outliers, and interpolate back to a dense scene flow field. The overview of the multi-frame pipeline is shown in Figure 2. A comparison between SFF and SFF++ is given in Table 2.

#### 4.1 Multi-frame Matching

*Initialization.* Improved initialization is the first extension in our multi-frame approach. We use the previous scene flow result and apply our assumption of constant motion. We propagate each 3D point according to its 3D motion to get an initial prediction of the scene flow at the current time step. We use the temporally propagated scene flow prediction as additional candidate during the initialization process earlier described in Section 3.1. The process is visualized in Figure 6. By not relying on the previous result alone, we successfully avoid error propagation (cf. Figure 6).

*Matching Cost.* As before, our matching process is based on the visual similarity of corresponding pixels (cf. Equation 1). We add the additional image correspondences to our scene flow matching cost from Equation 2 and




Fig. 6: Previous scene flow result (a), temporally warped scene flow (b), and multi-frame initialization (c) (on full resolution for visualization purposes).

obtain

$$\begin{aligned}
 C_{multi}(\mathbf{p}, \mathbf{u}) = & C(I_0^l, \mathbf{p}, I_1^l, \mathbf{p} + (u, v)^T) \\
 & + C(I_0^l, \mathbf{p}, I_0^r, \mathbf{p} + (-d_0, 0)^T) \\
 & + C(I_0^l, \mathbf{p}, I_1^r, \mathbf{p} + (u - d_1, v)^T) \\
 & + C(I_0^l, \mathbf{p}, I_{-1}^l, \mathbf{p} + (u_{-1}, v_{-1})^T) \\
 & + C(I_0^l, \mathbf{p}, I_{-1}^r, \mathbf{p} + (u_{-1} - d_{-1}, v_{-1})^T).
 \end{aligned} \tag{12}$$

We match all pixels of the image domain  $\Omega$  to all five other images and derive corresponding pixel locations from our scene flow representation in 2D (cf. Equation 12). For matching with the previous frame pair, we compute the inverse scene flow (in image space)  $u_{-1}$ ,  $v_{-1}$ ,  $d_{-1}$  by projecting the flow to 3D, inverting it, and projecting back to 2D, according to our constant motion assumption. Note that  $u_{-1} = -u$  if and only if  $d_0 = d_1$ . We do not simply invert the 2D optical flow directly. The (pixel-wise) constant motion assumption allows us to match across multiple time steps without increasing the search space or complexity of the scene flow domain.

*Visibility Prediction.* One additional major extension is our explicit visibility reasoning. Since the energy is based on visual data only, it is impossible to match regions, that are not visible in one of the images. Depending on the magnitude of camera movement, the baseline, and other circumstances, those unmatchable areas can become considerably large. Figure 7a shows an invisibility mask for pixels whose imaged 3D point is not visible in at least one of the views of  $I_0^r$ ,  $I_1^l$ , or  $I_1^r$ . This gives an impression of the limitations of dual-frame matching methods. However, Figure 7a also visualizes the remaining invisibility when considering one additional time step. More than two frame pairs can compensate for missing visual evidence when a pixel correspondence that is invisible in one image can be observed in another. Assuming to know which pixels



(a) Invisibility Example. Reference image (left), invisibility mask in the dual-frame scenario where no visibility handling is applied (middle), and remaining unmatched areas due to occlusions in the multi-frame scenario where visibility handling is applied (right).



(b) Visibility Prediction

Fig. 7: In the dual-frame case, occlusions can obscure large parts of the image which remain mostly visible when using multiple frames (a). Our explicit visibility handling of the multi-frame matching strategy predicts occluded and out-of-bounds regions (red) for all five images that are matched to the reference frame  $I_0^l$  (b).

are occluded or out-of-bounds of the image domain in each view, we can replace the matching cost for a single image correspondence (Equation 1) in Equation 12 by

$$C_{\text{vis}}(I, \mathbf{p}, I', \mathbf{p}', occ', oob') = \begin{cases} \theta_{occ}, & \text{if } occ'(\mathbf{p}) \\ \theta_{oob}, & \text{if } oob'(\mathbf{p}) \wedge \mathbf{p}' \notin \Omega \\ \theta_{penalty}, & \text{if } oob'(\mathbf{p}) \neq \mathbf{p}' \in \Omega \\ C(I, \mathbf{p}, I', \mathbf{p}'), & \text{else,} \end{cases} \quad (13)$$

with  $occ'$  and  $oob'$  being binary occlusion and out-of-bounds masks that indicate for each pixel  $\mathbf{p}$  whether the corresponding point in view  $I'$  is visible. For our full multi-frame matching cost in Equation 12, we need two invisibility masks for all 5 corresponding images. The adjusted cost uses a constant invisibility cost  $\theta_{occ}$  in case of occlusions to avoid trivial solutions where all points would be occluded. Same holds for out-of-bounds motions, with the additional option to assign a very large penalty cost  $\theta_{penalty}$  if motion under test is inconsistent with the out-of-bounds mask  $oob'$ . Otherwise, the normal

matching cost of Equation 1 will be used. In practice, we choose  $\theta_{occ} = \theta_{oob} = 10\,000$  based on empirical studies of the matching cost of ground truth scene flow vectors and make  $\theta_{penalty} = 10^6$  large enough so that out-of-bounds motion is forced to be consistent with the prediction.

Because the algorithm is designed to process sequences sequentially, we can use the estimated scene flow of the previous time step to predict the visibility for all six images. To this end, we use the temporally propagated scene flow prediction as for the initialization. This propagation can be used to check which parts of the scene leave the image domain  $\Omega$  for a specific view. These areas are marked as invisible in the associated out-of-bounds mask. It also allows to reason about occlusions by z-buffering, since we have full 3D information including depth. If multiple motions have the same target pixel in the target view, all but the closest are occluded. Examples of the visibility prediction are given in Figure 7b. For each of the five relevant frames, we show which pixel position of the reference frame can not be observed from the respective view. All pixels of the reference frame are visible by definition.

The impact of our multi-frame strategy combined with the explicit visibility reasoning will be shown in [Section 5.2](#) by comparing to the basic dual-frame SFF from [Section 3](#) ([Schuster et al., 2018c](#)).

*Consistency Check.* The next step of the pipeline is the consistency check to remove possible outliers. As before, we do so by computing a consistency scene flow field and compare corresponding scene flow vectors. Due to the changed setup to multiple frames from a stereo sequence, the way the consistency field is computed needs to be changed. We do not invert the temporal order of the images (cf. [Section 3.1](#)), but only change the view point. This leaves us with a left-right consistency check. Although a left-right check alone is less reliable than our previous consistency check, this has an important advantage. With inverted temporal order, scene flow correspondences according to the optical flow can not be established where the motion leaves the image boundaries. Thus, out-of-bounds regions are always filtered. This was no issue in the dual-frame approach, but using multi-frames we are able to match scene flow in these regions. With the left-right consistency check, we have the ability to maintain the correct correspondences in out-of-bounds regions. Matching for the consistency scene flow field is done exactly as before, using multiple scales, visibility reasoning, etc. Afterwards, each scene flow estimate of the reference frame is compared to its corresponding vector of the consistency field component-wise. If any error exceeds  $\tau_c = 1$  pixel, the whole scene flow vector is removed. SFF did benefit from an additional consistency check for the disparity only to obtain more matches for geometry interpolation. We follow that idea and compare our dense multi-frame matches to a second independent disparity map computed with the algorithm from [Yamaguchi et al. \(2014\)](#). Afterwards, we take the union of the additional matches and the results of the previous consistency check. We did change the stereo algorithm from SGM ([Hirschmüller, 2008](#)) to SPS ([Yamaguchi et al., 2014](#)) because it is faster, more accurate, and also able to match occluded and other invisible regions as our matching approach. The remaining merged correspondences are very accurate and serve as input for the robust interpolation (cf. [Figures 2](#) and [10](#)).

## 4.2 Robust Interpolation

We apply the same sparsification as in SFF because sparse-to-dense interpolation works best if input matches are not too dense already ([Bailer et al., 2015](#); [Hu et al., 2016](#); [Schuster et al., 2018c](#)). The sparsification works as an additional outlier filter and increases the spa-




Fig. 8: Sparsified input matches for interpolation (a) and the edge-aware neighborhood of seeds for the superpixel in white (b). Note that this is the actual size of segments on which we apply our interpolation models on. The supporting seeds are colored according to their distance to the superpixel.

tial support for the same number of neighbors during interpolation.

The interpolation in our extension is robust against possible remaining outliers and edge-preserving to create sharp motion boundaries at object edges. As in the less robust approach, we split interpolation into interpolation of the 3D geometry, followed by interpolation of the 3D motion. Because of this separation, we can use different sets of input seeds for each. Other than SFF in [Section 3](#) ([Schuster et al., 2018c](#)) where the interpolation algorithm of [Revaud et al. \(2015\)](#) was transferred to the higher dimensional scene flow problem, we adopt the robust concepts of [Hu et al. \(2017\)](#) for our scene flow problem.

In short, we dissect the reference frame into superpixels of size 25. Each superpixel is associated with a local neighborhood of the 200 closest input matches as shown in [Figure 8](#). The distance is computed as a geodesic distance based on an edge map. Additionally, each superpixel is initialized with two scene flow models, one for the 3D geometry and one for the 3D motion. Afterwards, the models for each superpixel is adjusted by propagation and random search. Based on the final models, dense scene flow can be computed for each pixel.

*Interpolation Models.* The geometric model is a slanted plane for each superpixel. It is initialized with a constant depth, i.e. parallel to the image plane. The motion model is a rigid transformation of 3D rotation and 3D translation for each superpixel. This rigid model is less universal, but more robust compared to our affine model in [Section 3.2](#). We do not assume rigid motions, though.

Our approach is not limited to rigid motions. Due to the small size of the superpixel segmentation, we can approximate non-rigid motions closely. Same holds for the planarity. Arbitrarily curved surfaces will be approximated by very small plane segments. The superpixel motion is initialized with translation only. For initialization, three different strategies have been tried: 1.) The basic approach of [Hu et al. \(2017\)](#) where the initial values are obtained from the the closest input seed. 2.) One where the local weighted medoid is used. 3.) One estimating the weighted geometric median. Though the multivariate estimators are much more accurate than the nearest neighbor method, the robust propagation with random search can compensate for less accurate initialization. Independent of the superpixel initialization, the final interpolation result was almost the same in all cases. We decided to use the robust geometric median in practice, since it provides a reasonable tradeoff between accuracy and run time.

*Robust Model Estimation.* To optimize the models for interpolation (plane + rigid transformation), we apply a robust, stochastic approach in a RANSAC-like fashion. For each superpixel, we randomly sample seeds from a local edge-aware neighborhood and use those to predict the models. Apart from random sampling, we use propagation to test estimated models at neighboring superpixels. To determine the fitness of a model, a truncated error for all supporting input seeds  $\mathbf{s}$  of the local neighborhood  $\mathcal{N}$  gets summed

$$C(M) = \sum_{\mathbf{s} \in \mathcal{N}} \min \left( \tau, \exp \left( -\frac{1}{\alpha} \cdot D(\mathbf{s}) \right) \cdot \epsilon(\mathbf{s}) \right). \quad (14)$$

In this formula,  $M$  is the respective model under test,  $\tau$  is the truncation error threshold of 4 pixels, and  $\epsilon(\mathbf{s})$  is the re-projection error when applying model  $M$  to input seed  $\mathbf{s}$  which is weighted by a geodesic distance  $D$  between the seed  $\mathbf{s}$  and the position of the respective superpixel with weight coefficient  $\alpha = 0.6$ . By minimizing the cost of all models for all superpixels using propagation and random model generation, we obtain robust interpolation models for geometry and motion. As before with multi-frame matching, our novel robust interpolation is compared directly to the previous interpolation method for scene flow matches of [Section 3.2](#) ([Schuster et al., 2018c](#)) in [Section 5.2](#).

*Universal Boundary Detector.* We will show that the boundary detector from [Section 3.2](#), trained on semantic boundaries of KITTI, performs much better than the original detector ([Dollár and Zitnick, 2013](#)) on that particular domain. However, improved robustness requires reliability across multiple domains. Thus, we train a




Fig. 9: Examples of boundaries from BSDS ([Martin et al., 2001](#)) (a), KITTI ([Geiger et al., 2012](#)) (b), and Sintel ([Butler et al., 2012](#)) (c) used to train our universal edge detector.

third variant with the goal to perform equally well on several different data sets. Towards this end, we use a joint set of training images from the 200 images of the original BSDS data set ([Martin et al., 2001](#)) and add semantic boundaries for 424 images of KITTI and for 100 images of Sintel. The semantics for KITTI are the same as before. For Sintel, we use the sequences *cave\_2* and *sleeping\_1* of the *clean* rendering pass that are excluded during evaluation. We create semantics by merging the provided mesh and material segmentation, e.g. all segments belonging to the dragon are labeled with the same ID. Examples of ground truth boundaries for each of the data sets are given in [Figure 9](#). By combining these three sets, we obtain a total of 724 images for training. The combined edge model is much more versatile and applicable to different data sets. Examples of the detected edges are visualized in [Figures 1a, 4, 8](#) and [12](#). Textures and shadows are suppressed, while at the same time, object boundaries are detected accurately.

For optional refinement of the dense 3D motion field, we use variational optimization of [Section 3.3](#) and the ego-motion model of [Section 3.4](#). However, our improvements in matching and interpolation make both refinement steps obsolete as we will show in [Section 5.1](#).

## 5 Experiments and Results

A series of experiments was conducted to evaluate our scene flow algorithm. We start with an ablation study to show the impact of each individual step of our sparse-to-dense pipeline and the incremental improvements presented within this paper. It follows a more detailed comparison between the dual-frame approach ([Section 3](#))

Table 3: Results of our scene flow estimation pipeline for different combinations of matching and interpolation settings, before and after variational optimization and ego-motion refinement.

Method	SF-bg	SF-fg	SF-all
dual+epic3d+bsds	26.93	32.16	27.73
dual+epic3d+bsds+var	26.14	30.71	26.84
dual+epic3d+bsds+var+ego	13.39	30.92	16.07
dual+epic3d+kitti	25.60	28.99	26.12
dual+epic3d+kitti+var	24.78	27.37	25.18
dual+epic3d+kitti+var+ego (SFF)	12.04	28.31	14.53
dual+ric3d+kitti	12.10	22.86	13.74
dual+ric3d+kitti+var	12.11	22.96	13.77
dual+ric3d+kitti+var+ego	11.57	23.46	13.40
multi+ric3d+kitti	10.75	19.37	12.07
multi+ric3d+kitti+var	11.26	19.60	12.54
multi+ric3d+kitti+var+ego	11.27	20.11	12.62
multi+ric3d+mixed (SFF++)	10.93	19.67	12.27
multi+ric3d+mixed+var	11.42	20.12	12.75
multi+ric3d+mixed+var+ego	11.29	20.62	12.72

and the robust multi-frame extension (Section 4). Afterwards, we will present results on the public KITTI scene flow benchmark and on the MPI Sintel dataset for comparison to state-of-the art. Finally, we disclose the limitations of our approach.

For all our experiments, we keep the parameters of SFF++ fixed at the given values in the previous sections. For SFF, we use the parameters of the original approach (Schuster et al., 2018c) which are not identical for KITTI and Sintel. The other approaches in our comparison also use adjusted parameters for the Sintel data set.

### 5.1 Ablation Study

The modularity of our method allows us to replace or leave out various parts of the pipeline. This way, it is easy to evaluate the effect of each component separately. We do so in Table 3 by evaluating the results of all 200 training images of the KITTI data set (Menze and Geiger, 2015) for several variants of our method. In particular, we vary the number of frames for the matching between *dual* and *multi*. We also alter the interpolation mechanism from the original one presented by Schuster et al. (2018c) (*epic3d*) to the robust version presented within this work in Section 4.2 (*ric3d*). Further, we evaluate the impact of the training data for the boundary detector. The original version was trained on BSDS (Martin et al., 2001) (*bsds*). Schuster et al. (2018c) presented a specialized version optimized for the KITTI data set (*kitti*). Lastly, we present results for our universal boundary detector (cf. Section 4.2) trained on a mix of data from BSDS, KITTI, and Sintel (Butler et al., 2012) (*mixed*). Finally, we successively add the varia-

tional refinement (*var*) and the ego-motion optimization (*ego*) for different combinations of the components.

In conformity with previous experiments (Schuster et al., 2018c), the dedicated KITTI boundary detector outperforms the original version that was trained on BSDS only. The universal boundary detector that was trained on mixed data performs almost equally well on KITTI and much better on other data sets like Sintel. Further, the novel more robust interpolation improves the accuracy of the scene flow result greatly, making the ego-motion model and even the variational refinement in almost all cases obsolete. Multi-frame matching improves the results even more. The impact of the multi-frame setup will be analyzed in Section 5.2 in more detail.

### 5.2 Evaluation of the Extensions

For visual comparison of our robust multi-frame approach SFF++ to SFF (Schuster et al., 2018c), we show optical flow and disparity at the next time step along with a scene flow error map for full matching, filtered matches, and dense interpolation in Figure 10. The error maps give correct estimates in green and outliers in blue to red according to the KITTI metric. Our novel multi-frame approach with explicit visibility reasoning is able to match occluded areas (e.g. next to the tree and traffic signs) and out-of-bounds regions (e.g. the front of the car) that are both not visible in some of the relevant frames (cf. Figure 7b). In addition, our robust interpolation can handle regions with almost no input seeds (e.g. the lower left part of the image). Table 4 compares the same in terms of average end-point-error and average outliers on all KITTI training sequences for all ground truth pixels (KITTI *occ* data) and occluded regions only (*occ* without *noc*). Our novel method is already very accurate during matching with a much smaller average end-point-error compared to SFF. Unlike SFF, our multi-frame approach is able to match almost 40 % of the invisible areas. Though overall filtered results after the consistency check appear to be worse, our novel approach is able to retain 9.8 % of the matches in occluded areas where SFF filters almost everything. That makes the spatial distribution of our matches preferable to that of the dual-frame approach (cf. Figure 10). Because of this and due to our robust interpolation, final results are much better with only about one third of scene flow outliers in occluded areas and less than half the percentage of overall outliers.

The direct comparison in these experiments shows that our novel multi-frame matching strategy with explicit visibility reasoning and robust interpolation is superior to the original SFF (Schuster et al., 2018c). Even

Table 4: Comparison of intermediate results for SFF (Schuster et al., 2018c) and our novel robust multi-frame extension SFF++ on the KITTI training data.

		All pixels								Occluded pixels only							
		EPE [px]			Outliers [%]				Density [%]		Outliers [%]				EPE [px]		
		D1	D2	F1	D1	D2	F1	SF	Density	SF	D1	D2	F1	SF	D1	D2	F1
Matching	SFF	7.2	11.3	38.9	12.6	29.2	32.8	39.8	100.0	100.0	22.1	92.1	95.5	99.5	34.3	50.1	195.7
	SFF++	<b>2.7</b>	<b>4.3</b>	<b>9.2</b>	<b>11.4</b>	<b>20.9</b>	<b>23.7</b>	<b>31.8</b>	100.0	100.0	<b>14.9</b>	<b>39.8</b>	<b>55.3</b>	<b>62.1</b>	<b>5.3</b>	<b>7.5</b>	<b>25.9</b>
Filtered	SFF	<b>0.8</b>	<b>0.9</b>	<b>1.1</b>	<b>2.2</b>	<b>2.7</b>	<b>2.2</b>	<b>4.2</b>	38.8	0.3	19.5	76.0	75.2	79.0	4.6	13.3	79.9
	SFF++	0.9	1.1	1.6	2.6	4.7	5.3	8.0	<b>41.6</b>	<b>9.8</b>	<b>2.5</b>	<b>12.9</b>	<b>23.0</b>	<b>26.2</b>	<b>1.1</b>	<b>2.4</b>	<b>9.3</b>
Interpolated	SFF	<b>1.2</b>	3.7	16.1	6.2	16.2	23.8	26.1	100.0	100.0	9.9	43.6	62.4	64.2	2.0	11.4	86.1
	SFF++	<b>1.2</b>	<b>1.9</b>	<b>4.7</b>	<b>5.2</b>	<b>8.8</b>	<b>9.7</b>	<b>12.7</b>	100.0	100.0	<b>7.1</b>	<b>16.0</b>	<b>22.7</b>	<b>24.8</b>	<b>1.7</b>	<b>3.3</b>	<b>15.3</b>

Table 5: Results on the KITTI Scene Flow Benchmark (Menze and Geiger, 2015).

Method	D1			D2			F1			SF			Run Time
	bg	fg	all	bg	fg	all	bg	fg	all	bg	fg	all	
ISF (Behl et al., 2017)	4.12	<b>6.17</b>	4.46	<b>4.88</b>	<b>11.34</b>	<b>5.95</b>	5.40	<b>10.29</b>	<b>6.22</b>	<b>6.58</b>	<b>15.63</b>	<b>8.08</b>	600 s
PRSM (Vogel et al., 2015)	<b>3.02</b>	10.52	<b>4.27</b>	5.13	15.11	6.79	<b>5.33</b>	13.40	6.68	6.61	20.79	8.97	300 s
OSF+TC (Neoral and Šochman, 2017)	4.11	9.64	5.03	5.18	15.12	6.84	5.76	13.31	7.02	7.08	20.03	9.23	3000 s
OSF18 (Menze et al., 2018)	4.11	11.12	5.28	5.01	17.28	7.06	5.38	17.61	7.41	6.68	24.59	9.66	390 s
SSF (Ren et al., 2017)	3.55	8.75	4.42	4.94	17.48	7.02	5.63	14.71	7.14	7.18	24.58	10.07	300 s
OSF (Menze and Geiger, 2015)	4.54	12.03	5.79	5.45	19.41	7.77	5.62	18.92	7.83	7.01	26.34	10.23	3000 s
SFF++ (ours)	4.27	12.38	5.62	7.31	18.12	9.11	10.63	17.48	11.77	12.44	25.33	14.59	78 s
FSF+MS (Taniai et al., 2017)	5.72	11.84	6.74	7.57	21.28	9.85	8.48	25.43	11.30	11.17	33.91	14.96	<b>2.7 s</b>
CSF (Lv et al., 2016)	4.57	13.04	5.98	7.92	20.76	10.06	10.40	25.78	12.96	12.21	33.21	15.71	80 s
SSF (Schuster et al., 2018c)	5.12	13.83	6.57	8.47	21.83	10.69	10.58	24.41	12.88	12.48	32.28	15.78	65 s
PRSF (Vogel et al., 2013)	4.74	13.74	6.24	11.14	20.47	12.69	11.73	24.33	13.83	13.49	31.22	16.44	150 s

Table 6: Results on MPI Sintel (Butler et al., 2012). We give average outliers (EPE &gt;3 px) for disparity and optical flow on each sequence separately and averaged over all sequences. We compare PRSM (Vogel et al., 2015), OSF (Menze and Geiger, 2015), FSF (Taniai et al., 2017), SFF (Schuster et al., 2018c), and our novel SFF++.

Sequence	Disparity					Optical Flow				
	PRSM	OSF	FSF	SFF	SFF++	PRSM	OSF	FSF	SFF	SFF++
Average	15.99	19.84	15.35	18.15	<b>13.60</b>	<b>13.70</b>	28.16	18.32	22.20	18.47
alley_1	7.43	5.28	5.92	8.81	<b>3.98</b>	<b>1.58</b>	7.33	2.11	3.95	2.11
alley_2	<b>0.79</b>	1.31	2.08	1.73	1.27	1.08	1.44	1.20	<b>0.87</b>	1.01
ambush_2	41.77	55.13	36.93	51.72	<b>31.56</b>	<b>51.33</b>	87.37	72.68	83.84	76.00
ambush_4	24.09	24.05	23.30	37.78	<b>22.25</b>	<b>41.99</b>	49.16	45.23	42.65	61.88
ambush_5	17.72	19.54	18.54	25.52	<b>13.48</b>	25.23	44.70	<b>24.82</b>	29.86	32.96
ambush_6	29.41	26.18	30.33	37.13	<b>23.17</b>	<b>41.98</b>	54.75	44.05	47.65	59.26
ambush_7	35.07	71.58	23.47	<b>16.34</b>	24.62	<b>3.35</b>	22.47	27.87	7.35	9.99
bamboo_1	<b>7.34</b>	9.71	9.67	14.53	10.80	<b>2.41</b>	4.04	4.11	4.15	3.44
bamboo_2	<b>17.06</b>	18.08	19.27	19.89	18.90	3.58	4.86	3.65	3.97	<b>3.57</b>
bandage_1	21.22	19.37	20.93	<b>16.42</b>	17.46	<b>3.30</b>	18.40	4.00	4.03	4.10
bandage_2	22.44	23.53	22.69	21.77	<b>16.80</b>	<b>4.06</b>	13.12	4.76	9.06	4.56
cave_4	<b>4.27</b>	5.86	6.22	6.20	4.93	16.32	33.94	14.62	<b>12.95</b>	18.16
market_2	<b>5.27</b>	6.61	6.81	6.71	6.26	<b>4.77</b>	10.08	5.17	6.09	5.51
market_5	15.38	13.67	<b>13.25</b>	26.66	14.13	28.38	29.58	<b>26.31</b>	28.87	32.56
market_6	<b>8.99</b>	10.29	10.63	14.53	10.18	<b>10.72</b>	16.39	13.13	16.69	13.91
mountain_1	0.42	0.78	0.23	0.15	<b>0.02</b>	<b>3.71</b>	88.60	17.05	89.57	10.84
shaman_2	25.49	28.27	24.77	<b>21.13</b>	23.94	<b>0.46</b>	1.67	0.56	4.31	1.80
shaman_3	33.92	52.22	<b>27.09</b>	35.37	29.02	1.75	11.45	<b>1.31</b>	8.51	5.53
sleeping_2	<b>1.74</b>	2.97	3.52	3.07	2.24	<b>0.00</b>	0.01	0.02	0.03	<b>0.00</b>
temple_2	<b>4.92</b>	5.54	5.96	6.98	4.95	<b>9.51</b>	10.52	9.66	12.57	12.05
temple_3	11.04	16.62	10.65	8.61	<b>5.76</b>	32.10	81.39	62.34	49.18	<b>28.64</b>



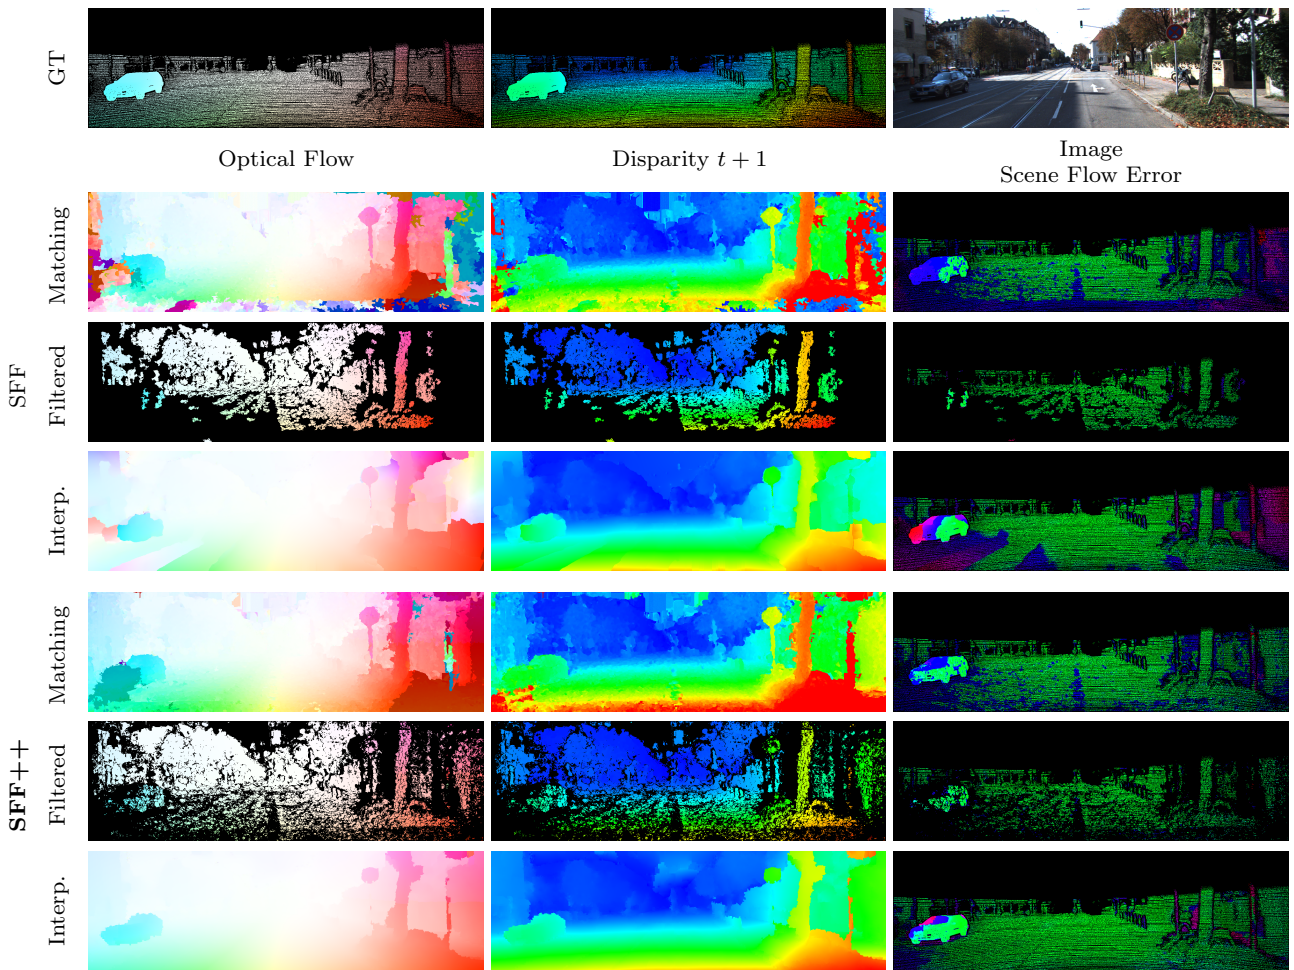


Fig. 10: Visual Comparison of SFF (Schuster et al., 2018c) and our novel SFF++. Optical flow, disparity at  $t + 1$  and scene flow error maps are shown for matching, filtering, and dense interpolation.

with the significant boost of the optional ego-motion model of Schuster et al. (2018c), our novel method without this model outperforms the dual-frame version on both evaluated data sets (see Sections 5.3 and 5.4).

### 5.3 KITTI Benchmark

We use the famous KITTI data set (Menze and Geiger, 2015) for the first part of our comparison to state-of-the-art. Please note that this benchmark does not provide reference motions for vulnerable road users like pedestrians or cyclists. The only dynamic objects are rigidly moving vehicles. Complete results for this benchmark are publicly available on the respective homepage. In Table 5 we give the top methods of the leader board, where the average amount of outliers are compared for disparity at both time steps ( $D1$ ,  $D2$ ), optical flow ( $FI$ ), and scene flow ( $SF$ ). Each category is further divided into regions of background ( $bg$ , static areas), foreground

( $fg$ , moving objects), and both ( $all$ , all available ground truth). As assumed, our novel approach scores better than the dual-frame conference version (Schuster et al., 2018c). Especially for the important foreground regions, our results are even comparable to the deep learning approach of Ren et al. (2017) that uses semantic segmentation and to that of Menze and Geiger (2015) that explicitly estimates a single rigid motion for complete independent objects. In addition, our method is at least three times faster than any better performing algorithm.

While investigating the discrepancy between our results on the validation data (Table 4) and on the test data (Table 5), we discovered that our multi-frame approach has difficulties in case the constant motion assumption is harshly violated (cf. Section 5.5). This can happen in KITTI due to the low frame rate of 10 frames per second and because of strong sudden pitch or roll rotations on bumpy streets, e.g. the last example in Figure 11. However, in most cases the assumption holds and produces robust and accurate results.




Fig. 11: Comparison of the scene flow error on the public KITTI Scene Flow Benchmark (Menze and Geiger, 2015) for PRSM (Vogel et al., 2015), OSF+TC (Menze et al., 2018), FSF+MS (Tani ai et al., 2017), SFF (Schuster et al., 2018c), and our novel approach SFF++.

We compare the error maps for different frames of the KITTI test data for PRSM, OSF+TC, FSF+MS, SFF, and our novel method SFF++ in Figure 11. Even for the partly occluded, badly illuminated vehicle in the second example, scene flow is estimated reliably. In this example also, scene flow at the occluded areas around the left traffic light is predicted correctly with sharp boundaries around the traffic sign. Within the third example, larger parts of the background exceed the error threshold slightly due to the violated constant motion assumption. However, dynamic objects are still detected more reliably than in most other approaches.

#### 5.4 MPI Sintel

The impact of our contributions has already been shown on a public benchmark in the previous section. To additionally demonstrate the robustness and versatility, we evaluate SFF and our more robust multi-frame extension SFF++ on a second data set – MPI Sintel (Butler et al., 2012). For our novel SFF++, we use exactly the same set of parameters. We explicitly highlight this, since many methods rely on parameter tuning. Espe-

cially deep learning models often require fine tuning. Sintel has a lot of contrary properties to KITTI. Most prominently, MPI Sintel consists of non-realistic, synthetically rendered images. Further, images are captured from a totally different domain and have therefore different characteristics. Sintel contains small as well as very large motion displacements of deformable, articulated characters. Thus, many of the included motions are non-rigid and the geometries are less often planar, compared to KITTI. This is also reflected in the results of this data set that are given in Table 6. OSF (Menze and Geiger, 2015), that is strongly relying on the piece-wise rigid plane model, performs considerably worse than on KITTI. Same is assumed for methods that rely on specialized deep neural networks, e.g. (Behl et al., 2017; Ren et al., 2017).

For evaluation, we have used the *final* rendering pass of all training sequences except *cave\_2* and *sleeping\_1* as Schuster et al. (2018c); Tani ai et al. (2017). Since no full scene flow ground truth is available, we compute the average percentage of outliers for disparity and optical flow, similar to *D1-all* and *F1-all* on KITTI. SFF++ outperforms the less robust dual-frame version

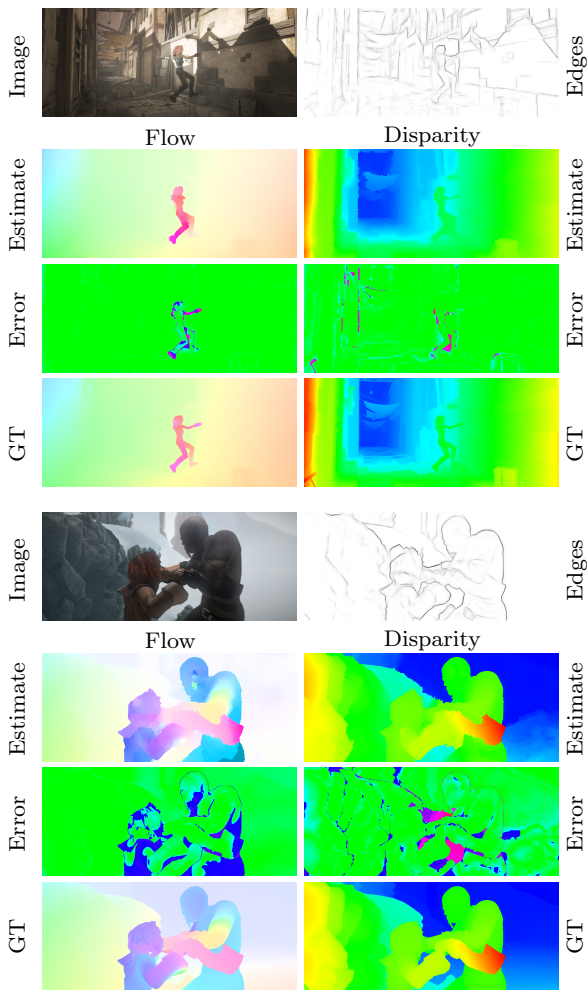


Fig. 12: Exemplary results of our method on the Sintel data set (Butler et al., 2012).

SFF (Schuster et al., 2018c), again highlighting the improvements of our contribution. We even achieve the best accuracy for disparity with quite some margin to the next best method FSF (Taniai et al., 2017) and the best performing method on KITTI, PRSM (Vogel et al., 2015). With that, our method joins the few methods with top performance on both data sets next to PRSM (Vogel et al., 2015), FSF+MS (Taniai et al., 2017), and SFF (Schuster et al., 2018c).

Figure 12 shows exemplary results of our approach for one frame of the sequences *alley\_2* and *ambush\_5*. The latter can be considered particularly challenging (cf. Table 6), though it is handled quite well by our approach. Figure 12 also demonstrates that a bit of accuracy is traded in for increased robustness. Small details, e.g. the arm of the girl, are lost during interpolation.

## 5.5 Limitations

Despite all efforts, it is not possible to avoid any assumption to estimate scene flow due to the ill-posedness of the problem. However, we tried to keep the method as unrestricted as possible. Our initial matching is purely data-based and requires no assumptions at all. The later imposed regularization during interpolation makes use of the piece-wise rigid plane model but applies it to an extreme over-segmentation of the scene into tiny superpixels of 25 pixels to reduce the negative effects greatly. Further, to make use of additional image information from multiple time steps, SFF++ assumes a constant motion within a temporal window of three stereo frames. With that, SFF++ is subject to a set of (mostly theoretical) limitations that will be described here.

As discussed several times, SFF++ keeps unaffected by the drawbacks of a piece-wise rigid, planar motion assumption even though this model is used during interpolation. The reason is the size of the superpixel segments. An impression of the size of the superpixels is given in Figure 8. During all our experiments, we could not observe that our superpixels are the limiting factor in the estimation of non-planar surfaces or non-rigid motions. However in theory, not all kinds of objects and motions can be approximated by small superpixels.

Another hypothetical failure case arises from the complete separation of matching and regularization. If unregularized matching leads to the removal of entire image regions during the consistency check, the later imposed regularization of the interpolation can not recover the content of these regions. This phenomenon was regularly observed for the dual-frame approach when highly dynamic objects were leaving the image domain. At the same time, this case was one of the motivations for the shift to multiple frames. In the multi-frame scenario, it is much less likely that matching fails consistently for entire regions. This is also supported by the study of visible areas in Figure 7a.

Lastly, the constant motion assumption is the only assumption-based limitation that causes practical impact on the performance of SFF++. For the KITTI data, we could notice degradation of the estimated scene flow due to the violation of this assumption. One of these examples is shown in the right column of Figure 11. Anyway, this problem was encountered rarely, mostly in the presence of potholes or crossing rails that lead to an extreme sudden discontinuity of the motion.

## 6 Conclusion

SFF is the first approach that uses sparse-to-dense interpolation for scene flow estimation. The extension to multi-frame matching with visibility prediction of SFF++ produces accurate correspondences even for occluded parts of the scene. A consistency check removes outliers. Robustly estimated models for 3D geometry and motion on small superpixels are used as novel interpolation method to reconstruct a dense scene flow field. Results on two diverse data sets that were computed with the exact same parameters have demonstrated the robustness for both, SFF and SFF++. Our robust multi-frame approach achieves the seventh best position in the overall ranking on KITTI and is among the top performing methods on MPI Sintel, with the lowest amount of outliers for disparity estimation. In a joint score on both data sets, we claim to achieve the second best performance after PRSM (Vogel et al., 2015). Separate fine tuning on each data set is assumed to further increase the accuracy of our method. However, the focus in this work was on robustness across different domains. To make our approach even more robust in the future, we suggest to use more than three frame pairs and most importantly evolve from our constant motion assumption to a higher order motion model.

## References

- Bailer C, Taetz B, Stricker D (2015) Flow Fields: Dense Correspondence Fields for Highly Accurate Large Displacement Optical Flow Estimation. In: International Conference on Computer Vision (ICCV)
- Bailer C, Taetz B, Stricker D (2018) Flow Fields: Dense Correspondence Fields for Highly Accurate Large Displacement Optical Flow Estimation. Transactions on Pattern Analysis and Machine Intelligence (PAMI)
- Basha T, Moses Y, Kiryati N (2013) Multi-view Scene Flow Estimation: A View Centered Variational Approach. International Journal of Computer Vision (IJCV)
- Behl A, Jafari OH, Mustikovela SK, Alhaija HA, Rother C, Geiger A (2017) Bounding Boxes, Segmentations and Object Coordinates: How Important is Recognition for 3D Scene Flow Estimation in Autonomous Driving Scenarios? In: Conference on Computer Vision and Pattern Recognition (CVPR)
- Brox T, Bruhn A, Papenberger N, Weickert J (2004) High Accuracy Optical Flow Estimation Based on a Theory for Warping. European Conference on Computer Vision (ECCV)
- Butler DJ, Wulff J, Stanley GB, Black MJ (2012) A Naturalistic Open Source Movie for Optical Flow Evaluation. In: European Conference on Computer Vision (ECCV)
- Chen Q, Koltun V (2016) Full Flow: Optical Flow Estimation by Global Optimization Over Regular Grids. In: Conference on Computer Vision and Pattern Recognition (CVPR)
- Dollár P, Zitnick CL (2013) Structured Forests for Fast Edge Detection. In: International Conference on Computer Vision (ICCV)
- Gaidon A, Wang Q, Cabon Y, Vig E (2016) Virtual Worlds as Proxy for Multi-Object Tracking Analysis. In: Conference on Computer Vision and Pattern Recognition (CVPR)
- Geiger A, Lenz P, Urtasun R (2012) Are We Ready for Autonomous Driving? The KITTI Vision Benchmark Suite. In: Conference on Computer Vision and Pattern Recognition (CVPR)
- He K, Sun J (2012) Computing Nearest-Neighbor Fields via Propagation-assisted kD-Trees. In: Conference on Computer Vision and Pattern Recognition (CVPR)
- Herbst E, Ren X, Fox D (2013) RGB-D Flow: Dense 3-D Motion Estimation Using Color and Depth. In: International Conference on Robotics and Automation (ICRA)
- Hirschmüller H (2008) Stereo Processing by Semiglobal Matching and Mutual Information. Transactions on Pattern Analysis and Machine Intelligence (PAMI)
- Hu Y, Song R, Li Y (2016) Efficient Coarse-to-fine Patchmatch for Large Displacement Optical Flow. In: Conference on Computer Vision and Pattern Recognition (CVPR)
- Hu Y, Li Y, Song R (2017) Robust Interpolation of Correspondences for Large Displacement Optical Flow. In: Conference on Computer Vision and Pattern Recognition (CVPR)
- Huguet F, Devernay F (2007) A Variational Method for Scene Flow Estimation from Stereo Sequences. In: International Conference on Computer Vision (ICCV)
- Jaimez M, Souiai M, Gonzalez-Jimenez J, Cremers D (2015) A Primal-Dual Framework for Real-Time Dense RGB-D Scene Flow. In: International Conference on Robotics and Automation (ICRA)
- Liu C, Yuen J, Torralba A (2011) SIFT Flow: Dense Correspondence Across Scenes and Its Applications. Transactions on Pattern Analysis and Machine Intelligence (PAMI)
- Lv F, Lian Q, Yang G, Lin G, Jialin Pan S, Duan L (2018) Domain Adaptive Semantic Segmentation Through Structure Enhancement. In: European Conference on Computer Vision (ECCV)
- Lv Z, Beall C, Alcantarilla PF, Li F, Kira Z, Dellaert F (2016) A Continuous Optimization Approach for Efficient and Accurate Scene Flow. In: European Con-

- ference on Computer Vision (ECCV)
- Ma WC, Wang S, Hu R, Xiong Y, Urtasun R (2019) Deep Rigid Instance Scene Flow. In: Conference on Computer Vision and Pattern Recognition (CVPR)
- Martin D, Fowlkes C, Tal D, Malik J (2001) A Database of Human Segmented Natural Images and its Application to Evaluating Segmentation Algorithms and Measuring Ecological Statistics. In: International Conference on Computer Vision (ICCV)
- Mayer N, Ilg E, Haussner P, Fischer P, Cremers D, Dosovitskiy A, Brox T (2016) A Large Dataset to Train Convolutional Networks for Disparity, Optical Flow, and Scene Flow Estimation. In: Conference on Computer Vision and Pattern Recognition (CVPR)
- Menze M, Geiger A (2015) Object Scene Flow for Autonomous Vehicles. In: Conference on Computer Vision and Pattern Recognition (CVPR)
- Menze M, Heipke C, Geiger A (2018) Object Scene Flow. *Journal of Photogrammetry and Remote Sensing (JPRS)*
- Neoral M, Šochman J (2017) Object Scene Flow with Temporal Consistency. In: Computer Vision Winter Workshop (CVWW)
- Ošep A, Hermans A, Engelmann F, Klostermann D, Mathias M, Leibe B (2016) Multi-Scale Object Candidates for Generic Object Tracking in Street Scenes. In: International Conference on Robotics and Automation (ICRA)
- Ren Z, Sun D, Kautz J, Sudderth EB (2017) Cascaded Scene Flow Prediction using Semantic Segmentation. In: International Conference on 3D Vision (3DV)
- Revaud J, Weinzaepfel P, Harchaoui Z, Schmid C (2015) EpicFlow: Edge-Preserving Interpolation of Correspondences for Optical Flow. In: Conference on Computer Vision and Pattern Recognition (CVPR)
- Ros G, Ramos S, Granados M, Bakhtiary A, Vazquez D, Lopez AM (2015) Vision-based Offline-Online Perception Paradigm for Autonomous Driving. In: Winter Conference on Applications of Computer Vision (WACV)
- Saxena R, Schuster R, Wasenmüller O, Stricker D (2019) PWOC-3D: Deep Occlusion-Aware End-to-End Scene Flow Estimation. In: Intelligent Vehicles Symposium (IV)
- Schuster R, Bailer C, Wasenmüller O, Stricker D (2018a) Combining Stereo Disparity and Optical Flow for Basic Scene Flow. In: Commercial Vehicle Technology Symposium (CVT)
- Schuster R, Bailer C, Wasenmüller O, Stricker D (2018b) FlowFields++: Accurate Optical Flow Correspondences Meet Robust Interpolation. In: International Conference on Image Processing (ICIP)
- Schuster R, Wasenmüller O, Kusch G, Bailer C, Stricker D (2018c) SceneFlowFields: Dense Interpolation of Sparse Scene Flow Correspondences. In: Winter Conference on Applications of Computer Vision (WACV)
- Taniai T, Sinha SN, Sato Y (2017) Fast Multi-frame Stereo Scene Flow with Motion Segmentation. In: Conference on Computer Vision and Pattern Recognition (CVPR)
- Vedula S, Baker S, Rander P, Collins R, Kanade T (1999) Three-Dimensional Scene Flow. In: International Conference on Computer Vision (ICCV)
- Vogel C, Schindler K, Roth S (2013) Piecewise Rigid Scene Flow. In: International Conference on Computer Vision (ICCV)
- Vogel C, Roth S, Schindler K (2014) View-consistent 3d scene flow estimation over multiple frames. In: European Conference on Computer Vision (ECCV)
- Vogel C, Schindler K, Roth S (2015) 3D Scene Flow Estimation with a Piecewise Rigid Scene Model. *International Journal of Computer Vision (IJCV)*
- Wang M, Deng W (2018) Deep visual domain adaptation: A survey. *Neurocomputing (Neurocomputing)*
- Wannenwetsch AS, Keuper M, Roth S (2017) ProbFlow: Joint Optical Flow and Uncertainty Estimation. In: International Conference on Computer Vision (ICCV)
- Wedel A, Rabe C, Vaudrey T, Brox T, Franke U, Cremers D (2008) Efficient Dense Scene Flow from Sparse or Dense Stereo Data. In: European Conference on Computer Vision (ECCV)
- Xu P, Davoine F, Bordes JB, Zhao H, Denœux T (2016) Multimodal Information Fusion for Urban Scene Understanding. *Machine Vision and Applications (MVA)*
- Yamaguchi K, McAllester D, Urtasun R (2014) Efficient Joint Segmentation, Occlusion Labeling, Stereo and Flow Estimation. In: European Conference on Computer Vision (ECCV)
- Yoshida T, Wasenmüller O, Stricker D (2017) Time-of-Flight Sensor Depth Enhancement for Automotive Exhaust Gas. In: International Conference on Image Processing (ICIP)
- Zweig S, Wolf L (2017) Interpo{N}et, a Brain Inspired Neural Network for Optical Flow Dense Interpolation. In: Conference on Computer Vision and Pattern Recognition (CVPR)