

Article

INV-Flow2PoseNet: Light-Resistant Rigid Object Pose from Optical Flow of RGB-D Images Using Images, Normals and Vertices

Torben Fetzter ^{1,*}, Gerd Reis ² and Didier Stricker ^{1,2}¹ Department of Computer Science, University of Kaiserslautern, 67663 Kaiserslautern, Germany² Department Augmented Vision, DFKI GmbH, 67663 Kaiserslautern, Germany

* Correspondence: torben.fetzter@dfki.de

Abstract: This paper presents a novel architecture for simultaneous estimation of highly accurate optical flows and rigid scene transformations for difficult scenarios where the brightness assumption is violated by strong shading changes. In the case of rotating objects or moving light sources, such as those encountered for driving cars in the dark, the scene appearance often changes significantly from one view to the next. Unfortunately, standard methods for calculating optical flows or poses are based on the expectation that the appearance of features in the scene remains constant between views. These methods may fail frequently in the investigated cases. The presented method fuses texture and geometry information by combining image, vertex and normal data to compute an illumination-invariant optical flow. By using a coarse-to-fine strategy, globally anchored optical flows are learned, reducing the impact of erroneous shading-based pseudo-correspondences. Based on the learned optical flows, a second architecture is proposed that predicts robust rigid transformations from the warped vertex and normal maps. Particular attention is paid to situations with strong rotations, which often cause such shading changes. Therefore, a 3-step procedure is proposed that profitably exploits correlations between the normals and vertices. The method has been evaluated on a newly created dataset containing both synthetic and real data with strong rotations and shading effects. These data represent the typical use case in 3D reconstruction, where the object often rotates in large steps between the partial reconstructions. Additionally, we apply the method to the well-known KITTI Odometry dataset. Even if, due to fulfillment of the brightness assumption, this is not the typical use case of the method, the applicability to standard situations and the relation to other methods is therefore established.

Keywords: light resistance; optical flow; scene flow; rigid alignment; point cloud registration

Citation: Fetzter, T.; Reis, G.; Stricker, D. INV-Flow2PoseNet: Light-Resistant Rigid Object Pose from Optical Flow of RGB-D Images Using Images, Normals and Vertices. *Sensors* **2022**, *22*, 8798. <https://doi.org/10.3390/s22228798>

Academic Editors: Jia-Rui Lin, Yu-Shen Liu, Xiaoping Zhou, Ge Gao, Anthony Tzes and Yi Fang

Received: 19 September 2022

Accepted: 9 November 2022

Published: 14 November 2022

Publisher's Note: MDPI stays neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Copyright: © 2022 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Three-dimensional reconstructions of objects and depth information of scenes play an increasingly important role in industry. Whether it is quality control in production or the recognition of the environment in autonomous driving, the number of applications is continuously increasing. Due to the simplicity of applicability, depth cameras are more and more used in parallel to flexible 3D scanners, and the availability of depth data for a wide variety of applications is steadily increasing. At the same time, the demand for scene understanding methods, represented by optical flow estimation, is constantly increasing, especially in the field of automation. Since in addition to images alone, more and more information is available, the demand for higher quality scene understanding also increases.

For the vast majority of applications, rigid scenes can be assumed and taken into account. In addition, even for dynamic scenes, the optical flow can be approximated by rigid models if not too large motions of the camera or the environment are expected. This rigidity assumption can even guide the estimation of optical flow, whose accuracy can benefit from it. The simultaneous extraction of the rigid transformation between two

subsequent frames is then also desirable. In this way, the method can be used for automatic alignment of point clouds in difficult scenarios, including large motion (fast driving cars) and large rotation (3D reconstruction, where often approx. 45° rotation between partial views occur) that yield strong shading changes.

The presented method uses an optical flow approach based on *PWC-Net* that has been adapted to use data from texture images, normal maps and vertex maps simultaneously. This optical flow method is moreover combined with the extraction of rigid transformations that are computed from the normal and vertex maps, warped by the predicted optical flow. The so predicted pose can mutually benefit from the coarse to fine strategy of the optical flow, since the optical flow can find dense correspondences over the whole scene using a pyramidal approach, even in the presence of large motion. Textural, geometric and shading features are included, which partly compensate for each other's weaknesses (sparsity of normal and vertex maps, illumination susceptibility of the texture images). From the warped 3D information of the scene, the rigid transformation can then be stably determined in a second step. Figure 1 depicts the basic methodology, which the presented approach will follow.

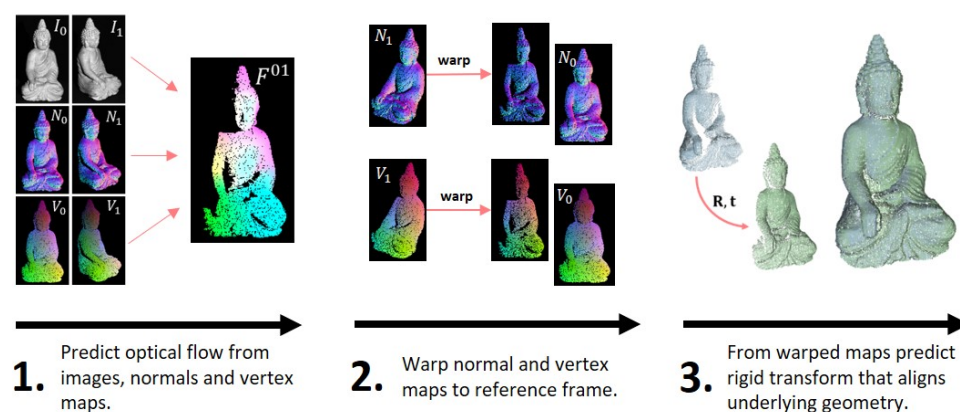


Figure 1. Sketch of the proposed methodology: In a first step, the pixel-wise optical flow is predicted from all available input (images, normals and vertices). In a second step, normal and vertex maps are warped to the reference frame using the predicted flow field. The stacked, warped normal and vertex maps are subsequently processed by another sub-network in order to predict a rigid transformation that aligns the underlying geometry.

Motivation: Flow-Based Alignment

In order to compute the alignment of two subsequently reconstructed frames, usually robust and transformation invariant features (SIFT, KAZE, ...) are detected and matched between the frames. Robust and outlier resistant methods such as RANSAC-based *PnP-solvers* are subsequently used to compute the rigid transformation between the views [1]. It is commonly known that this approach, applied with some few good features only, results in way better alignments than using many worse features jointly. Modern deep learning approaches adopt this scheme and deliver competitive results on a wide range of data in real time.

The basis of all common feature methods is the *brightness assumption*, which expects that the appearance of the object does not change significantly from one frame to another. This is fulfilled for many applications, especially when the camera moves smoothly through a scene or an object undergoes slow motion. If, on the contrary, the direction of the light incidence changes, the shading of the scene also differs dramatically and the brightness assumption gets strongly violated. This leads to a very probable failure of the standard methods based on this requirement, especially in the following situations:

- Outdoor scenes where lighting conditions can change suddenly. This can occur from direct sun light, as well as indirect light reflections from other objects.

- Moving objects, especially rotating ones, inevitably change the direction of light incidence. This leads in particular to considerable difficulties in the application area of 3D reconstruction, where the object is often rotated in order to capture it successively from all sides.
- Driving cars in the dark may cause strong shading differences in the captured images of the environment. Visible elements in the scene are illuminated by the car's headlights. These light sources move together with the car through the scene, which may yield strong variation of the direction of light incidence.

In order to illustrate this problem and investigate it, a setup with a static direct light source, a static camera and a rotated object is considered. Figure 2a shows how the standard approach based on SIFT matches fails, due to different light incidence. Figure 2b shows the scene in a different color coding that maps the grayscale values to a color scale that is more visual to human perception, which makes the different shading become obvious. While the features in the scene change in appearance, it can still be assumed that a significant portion of the scene overlaps in the different views. In the important case of object rotation in 3D reconstruction, our research shows that, in the vast majority of cases, a typical rotation of 45° still yields more than 80% overlap of the scene. Figure 2c visualizes the overlapping areas of the two views. Optical flow methods can benefit from this in turn, as they view and match the motion as a whole, using pyramidal approaches. Finally, Figure 2d shows correspondences determined using an optical flow method, as introduced in the following. The correspondences do contain noise and smaller errors, especially in feature-poor regions. They are nevertheless capable of predicting stable orientations of the object, significantly more stable ones than feature-based methods.




Figure 2. (a,d) show matches based on SIFT features and optical flow. (b) shows the scene, which has been illuminated by a strong spot light, in a different color space that is more visual to human perception. This more clearly visualizes the different shadings of the object, which is the reason for the failure of the common method based on SIFT features. (c) shows overlapping regions of subsequent scans. Even a rotation of approx. 45° yields a large overlap of more than 80%.

2. Related Work

Optical flow estimation is a well-known problem in applied machine vision and has wide spread use cases in industrial applications such as robotics, automotive driving, and quality control. The task is to determine dense motion at the pixel level between image pairs as accurately as possible. Starting with the method of Horn and Schunck [2], variational methods were the state of the art for a long time. Since the problem itself is an ill-posed problem, further assumptions have to be made on the flow field, which led to a multitude of different methods that use the most diverse regularization procedures to make the problem solvable according to the specific application. In recent years, the problem of *optical flow* estimation increasingly expanded to the problem of *scene flow* estimation, which deals with the 3D motion of scene points in space, whereas optical flow was limited to 2D point motion on the image plane. Based on the variational approaches for optical flow, a number of variational scene flow methods have been developed. Most of them use rectified stereo image pairs as input and thus estimate scene flow with different regularization methods or partial rigidity assumptions ([3–10]). At the same time, methods were developed to determine the scene flow directly from RGB-D data. With an increasingly number of depth

sensors that became available, this approach is quite justified. Several variants of methods handle this case ([11–14]).

The appearance of *FlowNet* [15] revolutionized the field of optical flow estimation. It became possible to treat the problem in real time with the help of *convolutional neural networks* (CNNs). In contrast, the variational methods were extremely time consuming and computationally expensive. A higher accuracy at the expense of a much larger network was subsequently achieved with *FlowNet2* [16]. This was followed by the release of *PWC-Net* [17], which uses warping layers at different levels of an image pyramid, representing the current state of the art that is in addition much smaller than the previously released *FlowNet2*. Based on *PWC-Net*, Saxena et al. have presented a method for estimating scene flow from rectified stereo image pairs. In addition, they handle occlusions within the forward pass. Previous methods required at least one forward and one backward warping to stably detect occlusions ([18–20]). Other approaches even tackle the task by iterative approaches such as [21]. In addition, a large amount of research currently focuses on either making networks lighter ([22,23]) or on training networks without ground truth through un- or self-supervision ([24–27]). A survey on variational as well as CNN-based optical flow methods can be found in [28].

Similar to earlier in the variational path, methods that extract scene flow directly from RGB-D data also evolved over time. Qiao et al. showed how scene flow based on *FlowNet* can be improved by fusion with features of depth data extracted in an extra network pass. Based on *PWC-Net*, Rishav et al. [29] use depth data from a Lidar sensor to determine scene flow. In doing so, they account for the lower resolution of Lidar data using appropriate reliability weights from [30]. In general, scene flow networks based on RGB-D data show poor performance for outdoor scenes, due to range limitations of the sensors. A number of approaches attempt to address this issue ([20,31–33]). Since the omission of active components removes the range limitations, but is accompanied by a loss of quality of the depth information, we will nevertheless restrict ourselves to this limited case. We are content with the scene flow within the sensor limits, since it is sufficient for an overwhelming number of practical applications, where the limits of the sensor can be planned accordingly. In order to predict the pose of an object, long time RANSAC approaches using explicit pose estimates based on the singular value decomposition were used. In recent years, first deep learning approaches predicted the pose directly using neural networks. Kendall et al. [34] use in their *PoseNet* several convolutional layers, followed by linear layers, to directly predict rotation and translation from RGB images. This way, they were the first to solve the problem of camera re-localization in static scenes by a deep learning approach. A few years later, Vijayanarasimhan et al. [35] extend this principle in *SfM-Net* in order to predict simultaneously the rigid transformations and the depth of the scene. They basically adopt the principles of the famous *Structure-from-Motion* pipeline to a deep learning framework. In parallel, Zhou et al. [36] developed a related model and showed how to train it in an un-supervised manner.

Finally, there has been a row of methods for direct point cloud registration with deep learning. Some of them replace parts of the standard strategies by deep learning methods and some try to replace the full pipeline. A large number of different approaches, correspondence-based and correspondence-free, are reviewed in [37,38].

Related to the presented work, Ref. [39] and recently Ref. [40] introduced variational and CNN-based methods for flow-aided pose estimation, based on fulfilled brightness assumption. Nevertheless, an automatic and light resistant flow-based pose-estimation method that works correspondence-free, and takes geometrical, textural and coherent scene motion into account has never been addressed before.

3. Light-Resistant Optical Flow

The *optical flow* between two images is understood as the displacements of the individual pixels from one to the other image. Determining the optical flow between images of a

scene often serves the purpose of scene understanding, as it directly allows the analysis of a large amount of scene information:

- The optical flow between calibrated camera images from different perspectives of the same static scene allows theoretically dense point correspondences and accompanying depth data.
- The optical flow between static camera images of a moving scene theoretically allows the analysis of scene motion and object tracking. If depth data are additionally available, the *scene flow*, i.e., the spatial movement of the points in the scene, can be calculated.

In the estimation of the optical flow between two consecutive images I_0 and I_1 , a horizontal and a vertical displacement field (F_x^{01} , F_y^{01}) are calculated, mapping each pixel in image I_0 to its corresponding pixel in image I_1 . The usual basis of the estimation is the brightness assumption, which assumes that corresponding pixels have the same appearance in the different images:

$$I_0(x, y) \approx I_1(x + F_x^{01}, y + F_y^{01}) \quad (1)$$

Figure 3 shows image I_0 and image I_1 , which has been warped by the optical flow F^{01} . Since the used optical flow has been computed from real data, the flow field is semi-dense and contains some masked pixels. Such errors will be addressed later on, where we will also show how to adopt filters to sparse, semi-sparse and mixed data. Instead of looking for exactly the same values between I_0 and I_1 , filtered values are considered in a regional context in order to robustify the matching. Deep neural networks have proven to be extremely effective for this purpose. The current state of the art is given by *PWC-Net*, which will be briefly introduced in the following to serve as a basis for the subsequently presented light-resistant method.




Figure 3. Image I_0 in comparison to image I_1 that has been warped by optical flow F^{01} . Assuming consistent brightness, these should be identical (ignoring masked pixels due to the semi-dense optical flow from real data). In case of strong rotations of the object, the shading changes dramatically, which violates this assumption.

3.1. PWC-Net

PWC-Net combines classical techniques such as a pyramidal approach, warping and correlation to create a highly effective network for optical flow estimation. The input images are passed through a pyramid of convolutions which extract rotation- and translation-invariant features at different levels of the receptive field. The number of hierarchies should be adapted that are appropriate to the image resolution. By successively halving the resolution in each step, the procedure should cover almost the entire scene in the filter of the last stage. From the lowest level, cost volumes based on extracted features are established from which the optical flow is effectively predicted. These flows are refined upwards with each level, incorporating new features of the current level and the flows and more global features from previous levels. By warping the data using the previous flow, the search space is significantly reduced and even large displacements can be treated and predicted with this comparatively small network. Figure 4 depicts the architecture of the network. Each prediction block consists of a cost volume for flow prediction and is fed with the corresponding layer in a U-Net structure, in order to predict a flow field in full resolution. Note that the standard network presented by Sun et al. in [17] predicts the optical flow

up to the second last level and afterwards refines the resulting flow by a context-network as post processing. This results in a final optical flow whose resolution is only $\frac{1}{16}$ of the input images' resolution. Instead of up-sampling by variational methods, we go for two additional texture-guided up-sampling steps within the network, in order to provide full resolution optical flows within a single training routine.

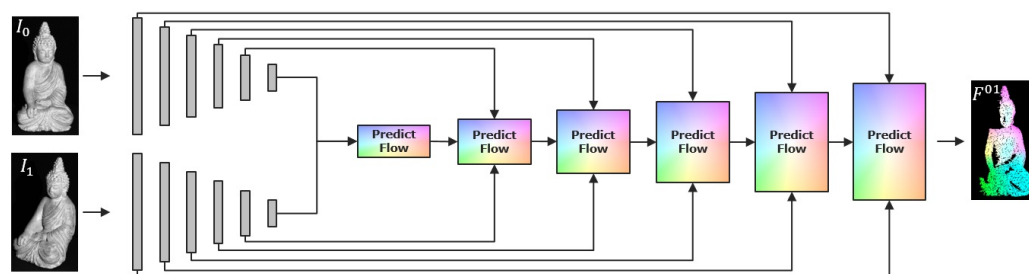


Figure 4. Sketch of the *PWC-Net* architecture. The input is convolved through multiple layers and the optical flow is predicted starting from the lowest level upwards in a U-Net structure. In each level, the layers of I_1 are warped towards the layers of I_0 in order to provide initial flows from previous lower levels. With this pyramidal approach, large flows are also predictable with quite small filter kernels.

3.2. INV-Net Using Images, Normals and Vertices

Classical *PWC-Net* uses texture images only. Unfortunately, for the investigated use case, these texture images may be disturbed due to shading changes, resulting from rotations of the object or light changes, which would make the network likely to fail due to a violated brightness assumption (see Figure 2). In many situations, where depth data are available, a lot of additional information can be provided to the network that is invariant under the shading effects related to light changes or object rotations:

- Texture images I_0 and I_1 that underlay shading effects. Nevertheless, they provide full and dense data, which can deliver local context.
- Depth maps D_0 and D_1 that store the relative geometrical information of the scene, light- and shading-invariant, with respect to the camera center. Due to measuring errors, there may be outliers or data-less pixels, resulting in semi-dense depth maps.
- Vertex maps V_0 and V_1 that store the spatial information of the scene, light- and shading-invariant in three channels of a map in image resolution. They are computed from the depth maps and the available camera calibration in order to store the geometrical information calibration independent. Therefore, they are similar to the depth maps semi-dense maps with masked pixels. Moreover, they are structured representations of point clouds that allow for performing neighboring operations on 3D data in 2D space, which yields large advantages in the following approach.
- Normal maps N_0 and N_1 that store spatial information of the surfaces in the scene. They are related to partial derivatives of the 3D vertices and do not underlay scaling and translation bias. They are in a specific range and responsible for a large amount of shading features of a scene (where standard methods based on fulfilled brightness assumption get a large amount of information from), without being disturbed by the light changes. They can be directly computed from the vertex maps, using the topological information given by the image grid (see [41]). Unfortunately, they thus also inherit the semi-density from underlying vertex maps.

Figure 5 sketches the basic problem of finding a light resistant pose estimation from all the available input. The first task is to find a light resistant high quality optical flow from this large amount of input data. In addition, both depth maps and vertex maps store the spatial information of reconstructed surface points. Since they are somehow interchangeable, we use the vertex maps only. This way, the method becomes independent of the intrinsic calibration to the cost of a higher amount of data that needs to be processed. Figure 9 (left part) sketches the basic network that takes features from *images* (textural features), *normal maps* (shading features) and *vertex maps* (geometrical features). Thereby, we

follow the basic principle of *PWC-Net* but run the different input through separate feature pipelines and set up independent cost volumes that contribute to the flow prediction. All features are processed as in [17] and fed to the pose prediction in each layer. This way, the network learns to treat the feature appropriate and to obtain advantages from all. Figure 6 depicts the prediction procedure in each layer, except the first one, where only the cost volumes are used for initialization of the flow.




Figure 5. Possible input that is available to the task of light resistant optical flow estimation and subsequent pose prediction. In addition to texture images, there are depth maps, vertex maps, point clouds and normal maps available. The depth maps as well as the vertex maps contain geometrical information. Since the vertex maps are independent of the calibration, it is the preferable choice for the presented method.

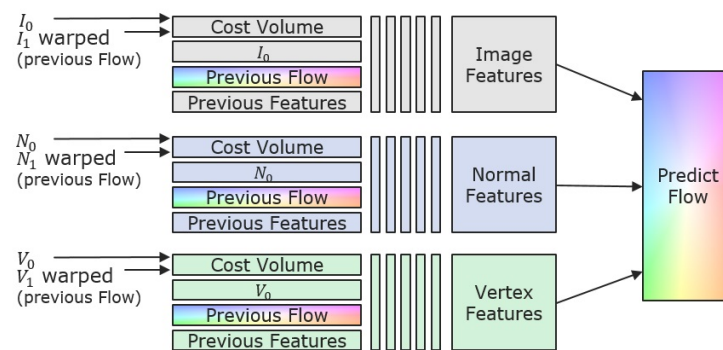


Figure 6. Flow prediction architecture in each layer (except first one). Features of images (texture), normals (shading) and vertices (geometry) are extracted separately and jointly fed to the prediction module.

3.2.1. Normalized Convolutions

In order to take into account the semi-density of the vertex maps and the normal maps, the convolutions, leading to the first layer, are replaced by *normalized convolutions* as introduced by Eldesokey et al. in [42]. Using the following slightly changed convolution procedure, the known masks can be used to ensure that data-less pixels do not contribute to the convolution with respect to neighbored pixels. Suppose, we are given a signal **A** to be convolved with a filter kernel **K**. Further assume that the measurements of the signal **A** are of varying quality with a confidence measure **W** of the same size as **A** having values between 0 and 1 to describe these uncertainties. It is desired to use the confidence measure as a weighting of the entries of **A** during convolution to ensure that reliable measurements have a higher influence on the convolution signal than inferior measurements or missing data for certain points. For this purpose, each summand within the convolution is weighted accordingly and divided by the sum of the weights to ensure the normalized character of

the convolution. In detail, the normalized convolution of signal \mathbf{A} , convolved with kernel \mathbf{K} and weighted by confidence \mathbf{W} around data point $[n]$, is given by

$$(\mathbf{K} * \mathbf{A})_{\mathbf{W}}[n] = \frac{\sum_m \mathbf{K}[m] \cdot \mathbf{W}[n-m] \cdot \mathbf{A}[n-m]}{\sum_m \mathbf{K}[m] \cdot \mathbf{W}[n-m]} = \frac{(\mathbf{K} * (\mathbf{W} \odot \mathbf{A})) [n]}{(\mathbf{K} * \mathbf{W}) [n]}, \quad (2)$$

where \odot denotes the element-wise *Hadamard-Product*. In order to avoid influence of missing pixels, a binary mask that contains zeros in case of missing data, and ones otherwise, can be fed to the convolutions as confidence \mathbf{W} .

3.2.2. Consistency Assumptions

Similar to the brightness assumption given in Equation (1), the following consistency assumptions hold true for normals and vertices of rigid scenes:

$$N_0(x, y) \approx \mathbf{R} N_1(x + F_x^{01}, y + F_y^{01}) \quad (3)$$

$$V_0(x, y) \approx \mathbf{R} V_1(x + F_x^{01}, y + F_y^{01}) + \mathbf{t} \quad (4)$$

Figure 7 visualizes the consistency relations for normals and vertices. While the pixels of the warped normal map coincide with the reference normals up to a rotation matrix \mathbf{R} , the vertices coincide up to rotation \mathbf{R} and a translation vector \mathbf{t} . These relations will be essential later on, in order to extract the rigid pose from the given optical flow. A very important result of our research is that features, computed from filtered normal and vertex maps, allow for computation of accurate optical flows. This means that the standard approach for feature extraction from images (as used in *PWC-Net*) is suitable to compute rotation- and transformation-invariant features from normal and vertex maps as well.

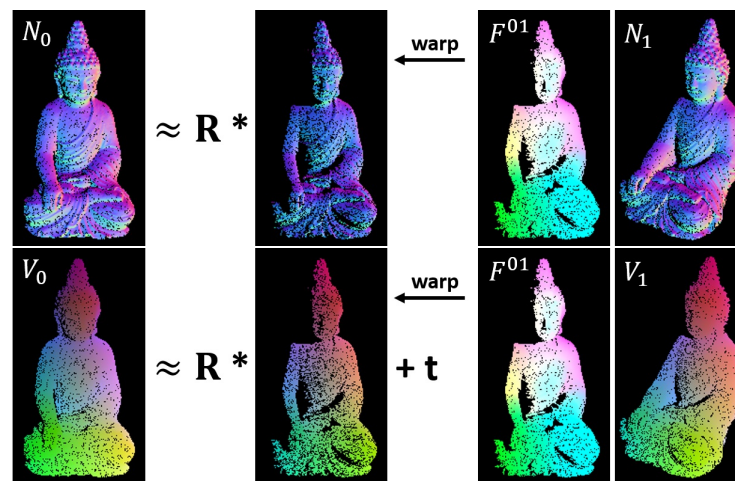


Figure 7. Normal maps and vertex maps that have been warped by optical flow F^{01} . Assuming rigid scenes, normals should be identical up to a rotation, vertices up to a rotation and a translation.

4. Pose from Warped Normals and Vertices

Several research works have already shown that it is possible to predict the relative pose of two views of a scene using neural networks. Usually, features are detected and matched, and outliers are rejected and then passed through a series of layers in order to obtain representative feature vectors. Finally, as introduced in [34], relative translation and rotation are predicted jointly using at least two subsequent fully connected layers.

In the previous section, a light-resistant optical flow has been computed by *INV-Net*. Based on this, it is not necessary to search for matches in the entire image. Considering images, normal maps and vertex maps from the first view and the ones from the second view that have been warped towards the first one with the computed optical flow, the data at each pixel-position theoretically match densely. Of course, there are also many erroneous

and inaccurate regions in the flow field, especially in feature-poor areas, where the flow is mainly interpolated. Previous work has shown that, in general, more accurate poses are estimated when only a few good features are used for the calculation, instead of many less good ones. This can also be achieved by feature extraction from the warped normal and vertex maps. It should be noted that, in areas where good features for the pose estimation can be found, good optical flows are also available. In a way, both the optical flow and the subsequently calculated pose are based on these same good features. Nevertheless, in the case of low quality features, as is the case with texture-poor and smooth surfaces, or even many false features due to light changes, we benefit from the more general information of the dense flow field.

In order to obtain best poses from the warped vertex and normal maps, we investigated two different approaches (*1 Step Method* and *2 Step Method*), and a combination *3 Step Method* that combines both approaches.

4.1. 1 Step Method

This approach uses the concatenated warped vertex maps to extract jointly rotation matrix \mathbf{R} and translation vector \mathbf{t} that align the vertex maps rigidly. The relation is based on consistency assumption (4). Note that, after warping, the matching vertices are theoretically placed at the same location in the concatenated input. Due to convolutional layers, the network is able to extract reliable locations, where a more accurate optical flow has been provided. The basic structure is shown in Figure 9 at branch 0. on the right.

4.2. 2 Step Method

This approach uses two steps to predict rotation and translation individually by two separate networks. Following the consistency property of Equation (3), the warped normal map N_1 and the reference normal map N_0 are related by a rotation matrix \mathbf{R} only. In a first step, this relative rotation \mathbf{R} is predicted by stacking N_0 and the warped N_1 and processing them through several convolutional layers, followed by two fully connected layers in order to predict optimal rotation with respect to the normals.

Based on the third consistency property of rigid transformations, given in Equation (4), the translation \mathbf{t} is predicted from the warped vertex map V_1 that has been rotated by matrix \mathbf{R} and the reference vertex map V_0 . Rotation matrix \mathbf{R} , from the previous step, has been applied in order to obtain dependency on the translation vector \mathbf{t} for this inference step only. The structure is again shown in Figure 9 at branches 1. and 2. on the right.

4.3. 3 Step Method

Rotation and translation are two fundamentally different operations that have a strong influence on each other. The smaller a rotation, the better it can be approximated linearly. Unfortunately, the joint extraction as in *1 Step Method* may yield inaccuracies in case of large rotations. In these situations, it may be beneficial to extract them separately like in the *2 Step Method*. Nevertheless, small rotational errors, from the first step of this approach, influence the predicted translation from the second step.

The idea of the *3 Step Method* is to first apply the *2 Step Method* to pre-align the vertex maps.

In a third step, a correctional rotation matrix $\tilde{\mathbf{R}}$ and a correctional translation vector $\tilde{\mathbf{t}}$ are jointly predicted from the warped and pre-transformed vertex map $\mathbf{R}V_1 + \mathbf{t}$ and reference vertex map V_0 . The final pose $P = (\hat{\mathbf{R}}, \hat{\mathbf{t}})$, as depicted in Figure 8, is then given by:

$$\hat{\mathbf{R}} = \tilde{\mathbf{R}}\mathbf{R}, \quad \hat{\mathbf{t}} = \tilde{\mathbf{R}}\mathbf{t} + \tilde{\mathbf{t}} \quad (5)$$

For extracting this correctional transformation, the *1 Step Method* is used. This is beneficial, since the correctional rotations are usually small, which makes it possible to predict the rotation and the translation jointly in order to avoid weaknesses of successive prediction as in the *2 Step Method*. The structure is again depicted in Figure 9 at branches 1.,

2. and 3. on the right. We investigated that the combined 3 Step Method performs best, as it compensates for the respective weaknesses of both methods.




Figure 8. Point clouds of the two exemplary views. The resulting transformation $P = (\hat{\mathbf{R}}, \hat{\mathbf{t}})$ aligns the point cloud of the first view to the one of the second view. The registered combined point cloud is shown besides.

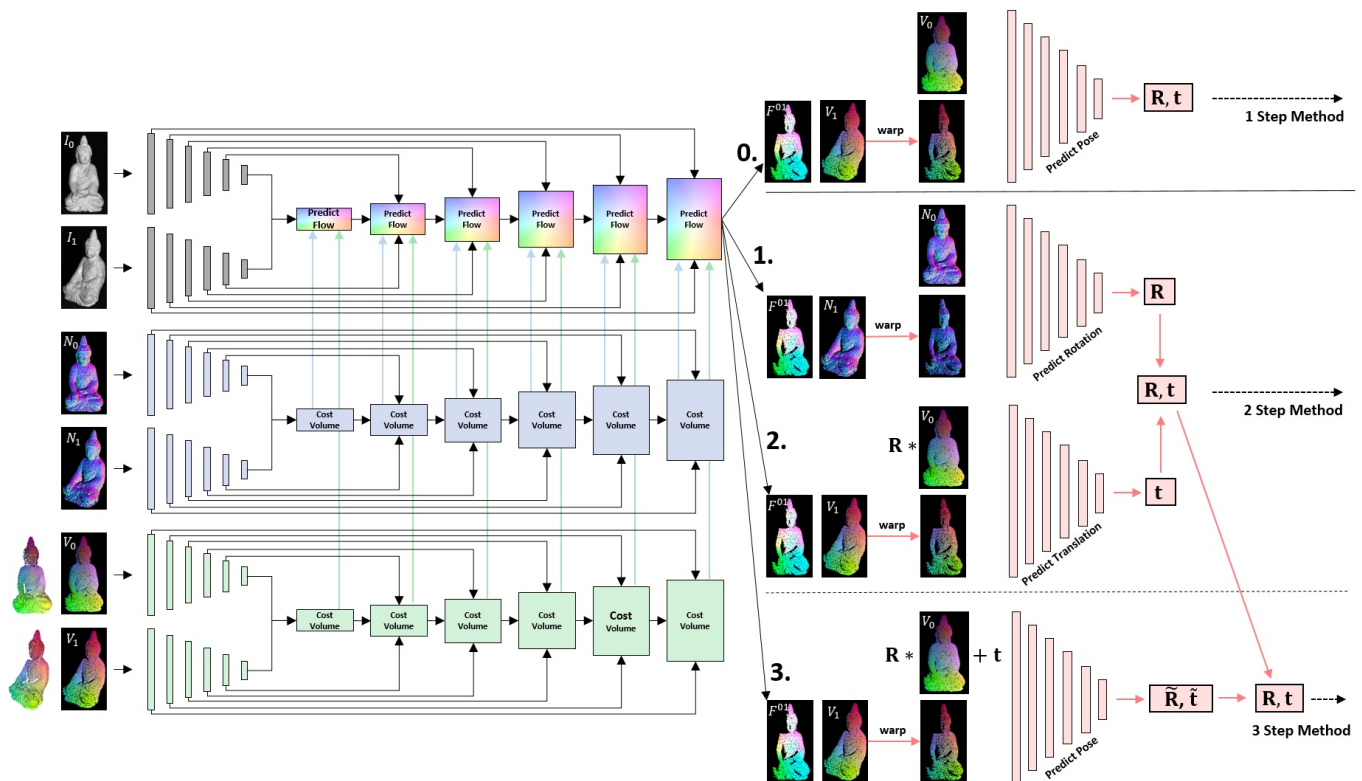


Figure 9. Architecture of Flow2PoseNet. The left part of the network aims to predict accurate optical flow from images, normal- and vertex-maps, using textural features from images, shading features from normals and geometrical features from vertices in order to predict accurate and light resistant flow fields. The pose of the rigid scene is computed in three steps from the warped normal- and vertex-maps. The first step predicts the normals from the warped normal-maps. The second step predicts the translation from the warped and rotated vertex-maps. The third step predicts a correction transformation to refine the predicted rotation and translation incrementally.

5. Data Sets and Data-Processing

There are already a number of public datasets, as well for optical flow estimation (*Flying Chairs*, *Sintel*, *Kitti*, *Flying Things3D*) as for pose estimation and odometry (*Kitty Odometry*, *3D Match*, *ModelNet14*, *ShapeNet*). Unfortunately, only datasets that provide both

images and depth data are suitable for the proposed investigations. Given the depth map and the camera calibration, the required normal maps can be approximated by practical methods, such as [41] and are thus not prerequisites. Therefore, for the evaluation of the estimated flow fields and the inferred poses in comparison to state-of-the-art techniques, the established *Kitty Odometry* dataset will be used later on. Even if it does not reflect the main application area for the development of the method, since it involves quite small rotations that barely show shading differences due to movement of the camera instead of the scene, it allows for comparison with previously existing methods.

Nevertheless, for the task of rotating objects, ground truth data of both optical flow and scene pose are required for training the presented network. In addition, it is advantageous to be able to use absolutely correct normal, depth and calibration data to avoid the influence of computational errors on the training. To the best of our knowledge, no such dataset exists. In addition, a general dataset for object orientation in the context of 3D reconstruction is not available to our knowledge. Therefore, several datasets are published together with this publication (<https://www.dfki.uni-kl.de/~fetzer/flow2pose.html> (accessed on 19 September 2022)). Among them are two synthetic datasets with rendered images, normals, depth maps and ground truths of camera calibration, optical flow and camera positions. One of them contains scenes with consistent scene illumination (*ConsistentLight*) of both camera views. The other one contains scenes with inconsistent illumination (*InConsistentLight*), where the position of the light source changes significantly between the views. This simulates the difficult case, where, for example, the object rotates, which may dramatically change the angle of incidence of the light (violated brightness assumption). The scenes of the synthetic data sets were created and rendered using *Unity* [43]. To avoid dependencies on the background, 75 spherical backgrounds were added to the scenes randomly. The grayscale images, depth maps, normal maps and optical flows were rendered for random scenes each from two random camera perspectives. The calibration information, the camera positions and the position of the illuminating point light are also provided. For both synthetic datasets, a training subset and a test subset were created. The training sets contain 20,000 random scenes in which objects were randomly placed in the scene. The test sets contain 1000 random scenes in which other objects that have not been used in the training sets were chosen. The 22 models used for the training sets are shown in Figure 10a and the eight models used for the test sets are shown in Figure 10b. Figure 11a–d shows the rendered data for an exemplary scene.

In a similar format, a real dataset (*BuddhaBirdRealData*) is delivered, which consists of captured data from five different objects, shown in Figure 10c. The images are captured by monochrome cameras. The depth data have been reconstructed by a structured light approach using a setup with a controlled environment. Thereby, the reconstructions have been performed within an approximately 1m³ working volume with a negligible ambient light component. Background effects were avoided by using the darkest possible background color. Camera as well as projector calibration information is provided along with the dataset. The normal maps are computed from the geometry, defined by the depth data and the calibration information. After manually aligning the data, the semi-dense flow fields have been computed and stored. The scenes were illuminated by a projector that has been calibrated jointly with the cameras and thus also delivers the light position in the scenes. Each model has been captured from eight positions with two different cameras each. Flow and pose data are available for each of the camera combinations of ancient positions, which yields ground truth data for 40 combinations per object. This results in 200 ground truth scenes of the real data that can be used for testing the models in real scenarios. Thereby, the first 40 pairs represent the scans within one scan head (consistent light) with eight reconstructions per object. The last 160 pairs represent the inconsistent light case with combinations of camera views between adjacent scans (that use different projectors). Similar to the synthetic case, Figure 11e–h shows the captured and estimated data for an exemplary real scene.

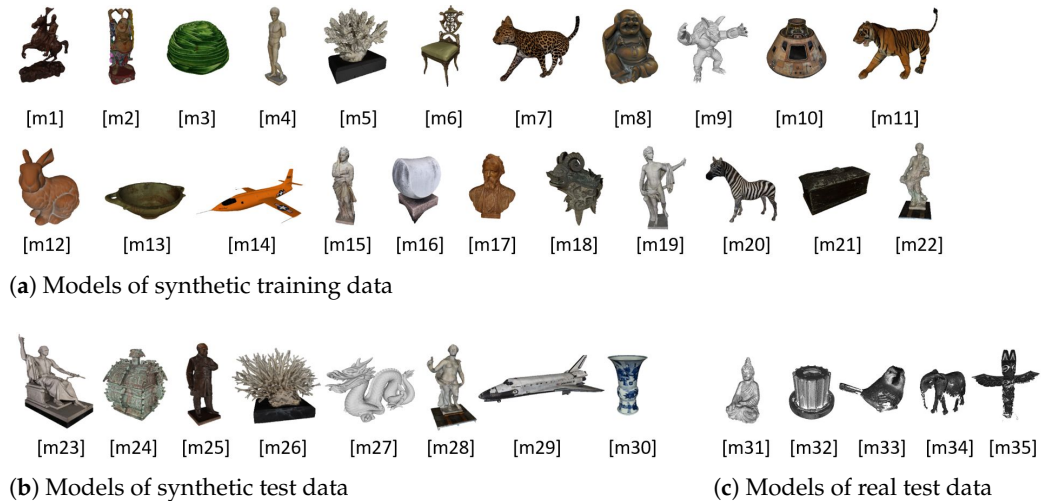


Figure 10. 3D models that have been used to create the synthetic and real datasets. (a) shows the models on which the synthetic training scenes are based on; (b) shows the models of the synthetic test scenes; (c) shows the models that result from the captured real data.

5.1. Data Sources and Data Formats

The 3D Models that have been used to create the data sets are taken from different sources and are free to use. Models [m9, m12, m27] were taken from the Stanford 3D Scanning Repository [44]. Models [m2, m7, m8, m11, m20] were taken from [45]. Models [m1, m3, m5, m6, m10, m13, m14, m17, m18, m21, m23, m24, m25, m26, m29, m30] were taken from the Smithsonian 3D Digitization page [46] that collected a large amount of 3D data from several museums and archives, from which many are free to use. Models [m4, m15, m16, m19, m22, m28, m31, m32, m33, m34, m35] resulted from our own research and are released with this work.

Each scene of the datasets, no matter if real or synthetic, consists of the following data parts:

- `image0` and `image1` contain the 8-bit integer grayscale images of the two camera views.
- `data0` and `data1` are .json files that contain the intrinsic calibration matrices \mathbf{K} , camera rotation \mathbf{R} and translation \mathbf{t} , the minimal and maximal depth values *minDepth* and *maxDepth*, the minimal and maximal values of the horizontal and vertical optical flows *minFlowX*, *maxFlowX*, *minFlowY* and *maxFlowY* and the coordinates of the light source *lightPos*.
- `depth0` and `depth1` are 16-bit integer grayscale images that need to be scaled after loading using minimal and maximal depth values from the data files:

$$D = D \cdot \frac{\text{maxDepth} - \text{minDepth}}{65535} + \text{minDepth}$$

- `norma10` and `norma11` are 24-bit integer RGB images in tangent space that can be re-transformed to spatial space by:

$$\mathbf{n} = \left(\frac{2\mathbf{n}_1}{255} - 1, \frac{2\mathbf{n}_2}{255} - 1, 1 - \frac{2\mathbf{n}_3}{255} \right)$$

- `flow0` and `flow1` contain the horizontal and vertical displacements of the respective flow fields between the views. The flows are stored as 16-bit integers in three channel images (*flowX*, *flowY*, *zeros*) and are scaled similar to the depth files.

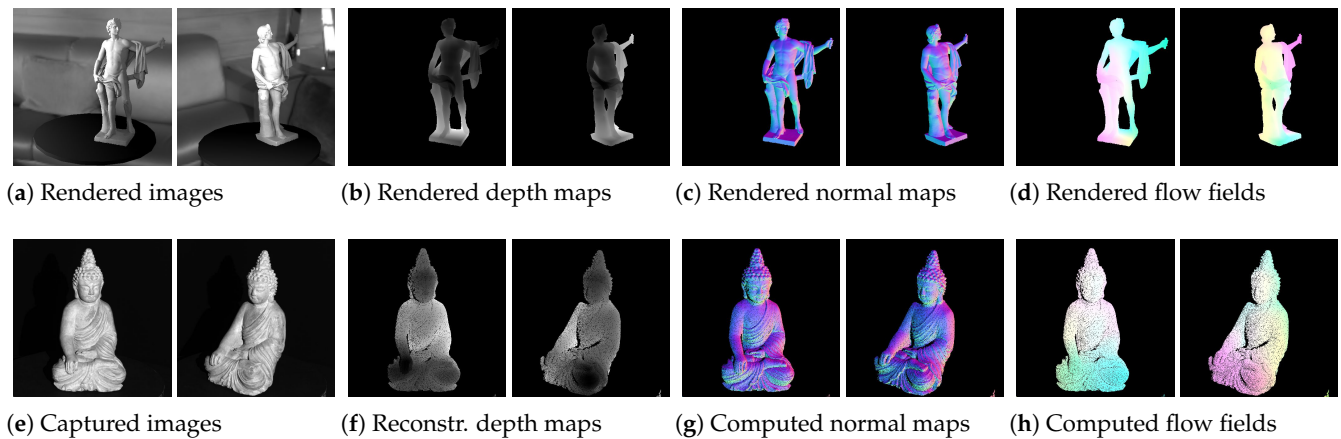


Figure 11. Example scene of the synthetic (top row) and real (bottom row) datasets. Each scene contains images, depth maps, normal maps and flow fields of two different camera views. In addition, a data file for each camera is stored that contains calibration information, camera position, light source position and minimal/maximal values of flows and depths in order to allow memory efficient saving of the data.

Note that missing/masked pixels for which no depth information is available contain zeros in the depth, flow and normal files. After re-scaling and shifting these files, the mask should be applied again to keep the masking information with values of zero.

The presented network uses vertex maps instead of depth maps. These can be computed from depth data and given calibration by applying the following operation to each image pixel (x, y) :

$$V(x, y) = \frac{\mathbf{K}^{-1}(x \ y \ 1)^T}{\|\mathbf{K}^{-1}(x \ y \ 1)^T\|_2} \cdot D(x, y) \quad (6)$$

5.2. Camera Pose and Scene Pose

The given depth, vertex and normal maps are independent of any camera pose, as these are usually not available beforehand and need to be computed by the procedure. In order to use them to triangulate point clouds with respect to the given pose, the vertex maps (or point clouds) and normal maps can be transformed in the following way. Given a camera pose $P = (\mathbf{R}, \mathbf{t})$, the 3D point with respect to a complete camera matrix $\mathbf{P} = \mathbf{K}[\mathbf{R}|\mathbf{t}]$ is given by:

$$V(x, y) = -\mathbf{R}^T \mathbf{t} + \mathbf{R}^T V(x, y) \quad (7)$$

and the normals of the respective 3D points are given by:

$$N(x, y) = \mathbf{R}^T N(x, y) \quad (8)$$

For completeness, remember that the camera itself is located in $\mathbf{R}^T \mathbf{t}$. In the usual case of unknown camera poses, only the relative transformation between two vertex maps/point clouds can be estimated from the given data by a procedure as introduced in the previous section. In order to use the provided data, to deliver relative ground truth transformations between two views, the absolute poses need to be transferred to relative ones. If we are given the camera extrinsics of two views $\mathbf{R}_0, \mathbf{t}_0$ and $\mathbf{R}_1, \mathbf{t}_1$, the relative pose between vertex map V_0 and vertex map V_1 is given by

$$\mathbf{R}_{01} = \mathbf{R}_1 \mathbf{R}_0^T, \quad \mathbf{t}_{01} = \mathbf{t}_1 - \mathbf{R}_1 \mathbf{R}_0^T \mathbf{t}_0 \quad (9)$$

where vertex map V_0 is mapped to vertex map V_1 by applying the transformation as:

$$V_1 = \mathbf{R}_{01} V_0 + \mathbf{t}_{01}. \quad (10)$$

Example code for reading, transforming and visualizing the data can be found with the datasets.

5.3. Pre- and Post-Processing of Data

Point clouds that need to be aligned may theoretically be of an arbitrary scale. Neural network based approaches, like the presented one, need to extract meaningful features within the given vertex maps to find corresponding points from which the desired transformation can be predicted. For this purpose, the network is adapted to the specific task with fixed weights that have been optimally determined during a data based training. For different absolute values of the 3D positions, it is not possible to extract meaningful features within the vertex maps with always the same weights. In particular, learned thresholds for activation within the network may not be applicable.

A practical way around this is to scale and move the point clouds, or equivalently the 3D data in the vertex maps, approximately towards the unit cube, which is located at the world origin. Within this working volume, the neural network can work effectively and perform the alignment. The calculated pose is then combined with the previous transformation towards the unit cube and thus provides the desired operation on the raw data.

In a first step, the point clouds are moved to the origin by subtracting the centroids. In a second step, the point clouds are scaled to fit approximately into the unit cube. Note that the method presented assumes the point clouds to be of similar scale, as it appears from usual depth data coming up from the same sensor. Therefore, the scaling factor s towards the unit cube should also be chosen similarly for both point clouds that are processed.

Let us be given the two point clouds $X_0 = \{x_1^{(0)}, \dots, x_M^{(0)}\}$ and $X_1 = \{x_1^{(1)}, \dots, x_N^{(1)}\}$ that need to be aligned. The centered point clouds at the origin are given by:

$$\begin{aligned} X_0 - \mu_0 &= \{x_n^{(0)} - \mu_0 \mid x_n^{(0)} \in X_0\}, \quad \mu_0 = \sum_{m=1}^M x_m^{(0)} \\ X_1 - \mu_1 &= \{x_m^{(1)} - \mu_1 \mid x_m^{(1)} \in X_1\}, \quad \mu_1 = \sum_{n=1}^N x_n^{(1)} \end{aligned} \quad (11)$$

$X_0 - \mu_0$ and $X_1 - \mu_1$ are then scaled jointly and robustly in order to ensure that 90% of the point clouds map into the according subspace of the unit cube ($[-0.45, 0.45]^3 \subset \mathbb{R}^3$) that is located at the origin. This robustifies the scaling and reduces the negative effect of outliers dramatically. Note that, in general, it can be assumed that at least 90% of a point cloud should contain usable data. Let us be given the set of values with maximal absolute coordinates of both centered point sets, $Y = \{\max(|x|) \mid x \in (X_0 - \mu_0) \cup (X_1 - \mu_1)\}$.

Having sorted the values $y_n \in Y$ in ascending order $y_1 \leq \dots \leq y_{M+N}$, the scaling factor that ensures 90% of both point clouds being mapped into the cube, defined above, is given by $s = 1/y_{\lfloor 0.45(M+N) \rfloor}$, where $\lfloor \cdot \rfloor$ denotes floor rounding to integer values. The scaled, centered point clouds that are ready to be fed to the network are finally given by:

$$\tilde{X}_0 = s(X_0 - \mu_0), \quad \tilde{X}_1 = s(X_1 - \mu_1) \quad (12)$$

Having computed a pose $\tilde{P} = (\tilde{\mathbf{R}}, \tilde{\mathbf{t}})$ using the neural network that aligns the scaled point clouds by

$$\tilde{\mathbf{R}} \tilde{X}_0 + \tilde{\mathbf{t}} \approx \tilde{X}_1, \quad (13)$$

the final transformation $P = (\mathbf{R}, \mathbf{t})$ that aligns the raw point clouds X_0 and X_1 is given by

$$\mathbf{R} = \tilde{\mathbf{R}}, \quad \mathbf{t} = \frac{1}{s}\tilde{\mathbf{t}} + \mu_1 - \tilde{\mathbf{R}}\mu_0 \quad (14)$$

6. Coherent Learning of INV-Flow2PoseNet

The goal is to train the network to estimate the best possible optical flow that will enable stable extraction of the pose. Therefore, to obtain an end-to-end trainable network, we define a joint loss function that penalizes both the ground truth flow and the extracted pose under given flow.

The *PWC-Net* structure predicts flows $F^{(l)}$ of different levels $l = 0, \dots, L$. The *Flow2PoseNet* moreover uses these flows in order to predict the relative rotation \mathbf{R} and translation \mathbf{t} . Let be given according ground truth $F_{\text{GT}}^{(l)}$, \mathbf{R}_{GT} and \mathbf{t}_{GT} .

6.1. Multiscale Endpoint Error

The multiscale endpoint error (EPE) penalizes the different levels of the flow calculation with different hardness, provided by the respective weighting parameters α_l :

$$\mathcal{L}_{\text{EPE}}(F^{(0)}, \dots, F^{(L)}) = \sum_{l=0}^L \alpha_l \|F^{(l)} - F_{\text{GT}}^{(l)}\|_{\mathcal{F}} \quad (15)$$

with suitable level weights α_l , $l = 0, \dots, L$ and Frobenius matrix norm $\|\cdot\|_{\mathcal{F}}$. In case of sparse data, the differences inside the norm are masked in order to take the sparsity into account.

Note that the higher levels, which describe the rather coarse flow, are more important than the lower levels, which obtain the higher levels as input. However, since the higher levels have a lower resolution, the flow errors in absolute numbers are smaller than those of the lower levels. As a rule of thumb, because of the pooling between each level, the weighting should be at least halved each time to account for the resolution discrepancy. The weights that have been used for the proposed network are $\{\alpha_0, \dots, \alpha_6\} = \{0.001, 0.0025, 0.005, 0.01, 0.02, 0.08, 0.32\}$.

6.2. Alignment Error

A measure that treats both rotation and translation together is the well-known alignment error. It models the mean Euclidean distance of all point correspondences given by the groundtruth flow:

$$\mathcal{L}_{\text{AE}}(\mathbf{R}, \mathbf{t}) = \sum \| \mathbf{R}V_0(x, y) + \mathbf{t} - V_1(x + F_x^{01}, y + F_y^{01}) \|_{\mathcal{F}} \quad (16)$$

This measure best describes the problem to be solved. It has the advantage that it weights the impact of rotation against the translation. Note that it is important to mask errors that contain invalid pixels either of V_0 or of warped V_1 , in order to ensure that only locations are taken into account, where matching vertices in both views are available.

Note that this error alone might erroneously interchange rotations and translation effects in order to receive a minimal alignment error. These interchanges can be prevented by adding some direct translational and rotational error terms to the overall loss function. These additional terms act as a regularization to enforce a better decomposition into translation and rotation.

6.3. Translational and Rotational Errors

The error of the predicted translation is given directly as the Euclidean distance towards the ground truth translation:

$$\mathcal{L}_{\text{TRANS}}(\mathbf{t}) = \|\mathbf{t} - \mathbf{t}_{\text{GT}}\|_2 \quad (17)$$

Special attention is required for the rotation error. A suitable differentiable error between two rotation matrices \mathbf{R} and \mathbf{R}_{GT} is given by the angular error, which is defined by the absolute value of the rotation angle θ of the relative rotation $\mathbf{R}_{rel} = \mathbf{R}\mathbf{R}_{GT}^T$. Having a look at the conversion towards the axis angle representation, there are basically two ways to compute the rotation angle. The first relation is given with the trace of the rotation matrix:

$$\text{Tr}(\mathbf{R}_{rel}) = 1 + 2 \cos(\theta) \quad (18)$$

Another way is to calculate the rotation angle from the length of the extracted rotation axis. Having an explicit rotation matrix, the rotation axis \mathbf{u} is given by:

$$\mathbf{R}_{rel} = \begin{pmatrix} a & b & c \\ d & e & f \\ g & h & i \end{pmatrix} \Rightarrow \mathbf{u} = \begin{pmatrix} h-f \\ c-g \\ d-b \end{pmatrix} \quad (19)$$

The rotation angle θ is related to the length of \mathbf{u} by:

$$\|\mathbf{u}\|_2 = 2 \sin(\theta) \quad (20)$$

A direct computation of θ from one of equations (18) or (20) requires the use of one of the inverse trigonometric functions *arcus sinus* or *arcus cosinus*. These yield numeric problems due to singularities in case of angles close to $\pm \frac{\pi}{2}$ or $\pm \pi$, which is unsuitable for a general loss function that has to be differentiable. A more stable way to achieve θ is to use the two-dimensional *arcus tangens* *atan2* with both arguments:

$$\mathcal{L}_{ROT}(\mathbf{R}) = |\text{atan2}(\|\mathbf{u}\|_2, 1 - \text{Tr}(\mathbf{R}_{rel}))| \quad (21)$$

6.4. Joint Training Loss

The joint loss function is subsequently given by:

$$\begin{aligned} \mathcal{L} = & \mathcal{L}_{EPE}(F^{(0)}, \dots, F^{(L)}) + \mathcal{L}_{AE}(\mathbf{R}_{1Step}, \mathbf{t}_{1Step}) \\ & + \mathcal{L}_{AE}(\mathbf{R}_{2Step}, \mathbf{t}_{2Step}) + \mathcal{L}_{AE}(\mathbf{R}_{3Step}, \mathbf{t}_{3Step}) \\ & + \mathcal{L}_{TRANS}(\mathbf{t}_{3Step}) + \mathcal{L}_{ROT}(\mathbf{R}_{3Step}) \end{aligned} \quad (22)$$

At the beginning of the training, the gradients of the computed optical flow are detached before back-propagating the alignment errors.

6.5. Representation of Rotation

In order to ensure the predicted rotation matrix to be a proper rotation, a minimal parameterization by *Euler Angles* is chosen. Therefore, three values (θ, ρ, ϕ) are predicted by the network, defining the rotation angles around the x , y and z axes by the rotation matrices:

$$\mathbf{R}_x = \begin{pmatrix} 1 & 0 & 0 \\ 0 & \cos(\theta) & -\sin(\theta) \\ 0 & \sin(\theta) & \cos(\theta) \end{pmatrix}, \mathbf{R}_y = \begin{pmatrix} \cos(\rho) & 0 & \sin(\rho) \\ 0 & 1 & 0 \\ -\sin(\rho) & 0 & \cos(\rho) \end{pmatrix}, \mathbf{R}_z = \begin{pmatrix} \cos(\phi) & -\sin(\phi) & 0 \\ \sin(\phi) & \cos(\phi) & 0 \\ 0 & 0 & 1 \end{pmatrix} \quad (23)$$

The total rotation is given by the consecutive execution of these rotations: $\mathbf{R} = \mathbf{R}_x \mathbf{R}_y \mathbf{R}_z$. Vice versa, the respective *Euler Angles* can be extracted from a given rotation matrix \mathbf{R} by:

$$\theta = \text{atan2}(-\mathbf{R}_{23}, \mathbf{R}_{33}) \quad (24)$$

$$\rho = \text{atan2}(\mathbf{R}_{13}, \sqrt{\mathbf{R}_{23}^2 + \mathbf{R}_{33}^2}) \quad (25)$$

$$\phi = \text{atan2}(-\mathbf{R}_{12}, \mathbf{R}_{11}) \quad (26)$$

This conversion is especially used to compute the *Euler Angles* of the refined rotation matrix $\hat{\mathbf{R}}$ in the *3 Step Method* of Section 4.

7. Evaluation

For evaluation, we compare the calculated optical flows and registrations qualitatively on different synthetic and real data sets. Highly accurate results visualize a good generalization without fine-tuning from synthetic training data to the difficult real test scenes.

Therefore, multiple positions of the real scene are shown in Figure 12. It visualizes the performance of the method applied to eight partial scans of the Buddha scene (from the *BuddhaBirdReal* dataset), as it usually comes up from 3D scanners. Using the alignment given by the neural network, a few iterations of *Iterative Closest Points (ICP)* for refinement yield impressive results on the overall aligned point cloud of the object.

Figures 13 and 14 show the results for exemplary objects from the training (top 3 rows) and test datasets (bottom 3 rows) for the consistent and inconsistent light (moving light source) case. Thereby, the first columns show the input data consisting of images, normals and depth maps (that are converted to vertex maps using the calibration information, as in Equation (6)). The second column shows the resulting optical flow in comparison to the semi-dense ground truth optical flow in column 3. Columns 4 and 5 finally show the initial and the registered point clouds using the proposed neural networks. Special attention should be given to row 6 of Figure 14, which shows the performance of the neural network on a real test scene without fine-tuning.

In particular, for comparability with other methods, we also consider a network trained on the popular training sequences of Kitti Odometry and evaluate it on the test data as shown in Figure 15. As the Kitti dataset has less strong rotations and less shading changes, it is not the typical use case for the proposed method. Nevertheless, the proposed method works reliably for these easier kind of situations as well.

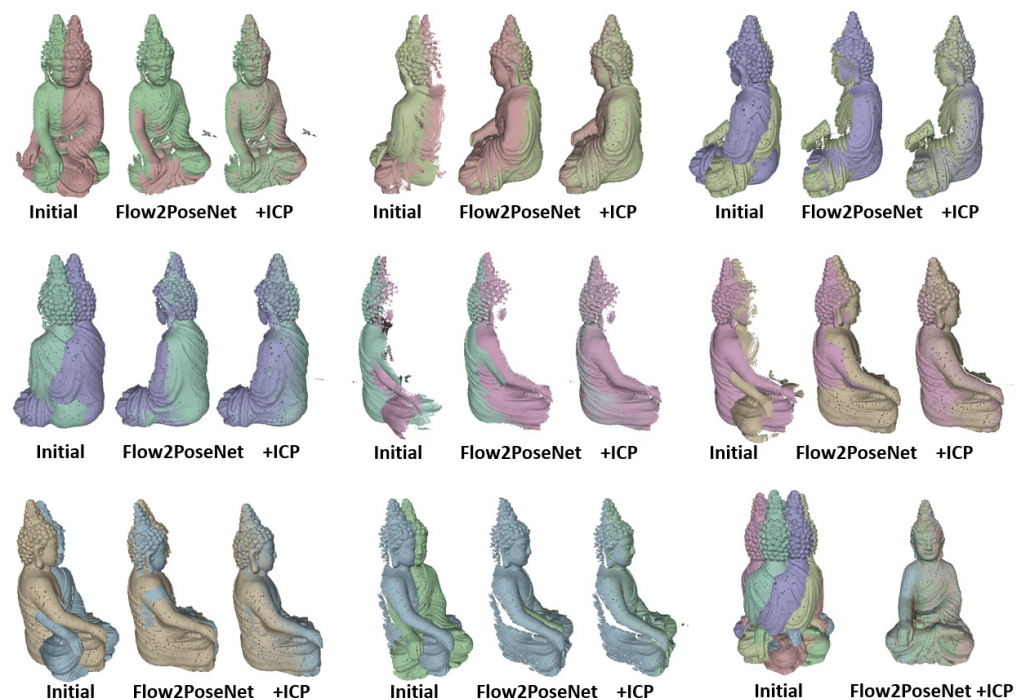


Figure 12. Application of the method to a full sequence of partial reconstructions of a real Buddha object from the *BuddhaBirdReal* dataset. Such sequences usually result from 3D scanners (as here from a structured light scanner). Since often a turntable is used, strong rotations ($\approx 45^\circ$) and shading changes disturb the data. After pre-alignment, a few iterations of the *ICP* algorithm are applied to refine the alignment of the point clouds. The image on the bottom right shows the impressive result on the overall aligned full point cloud of the statue.

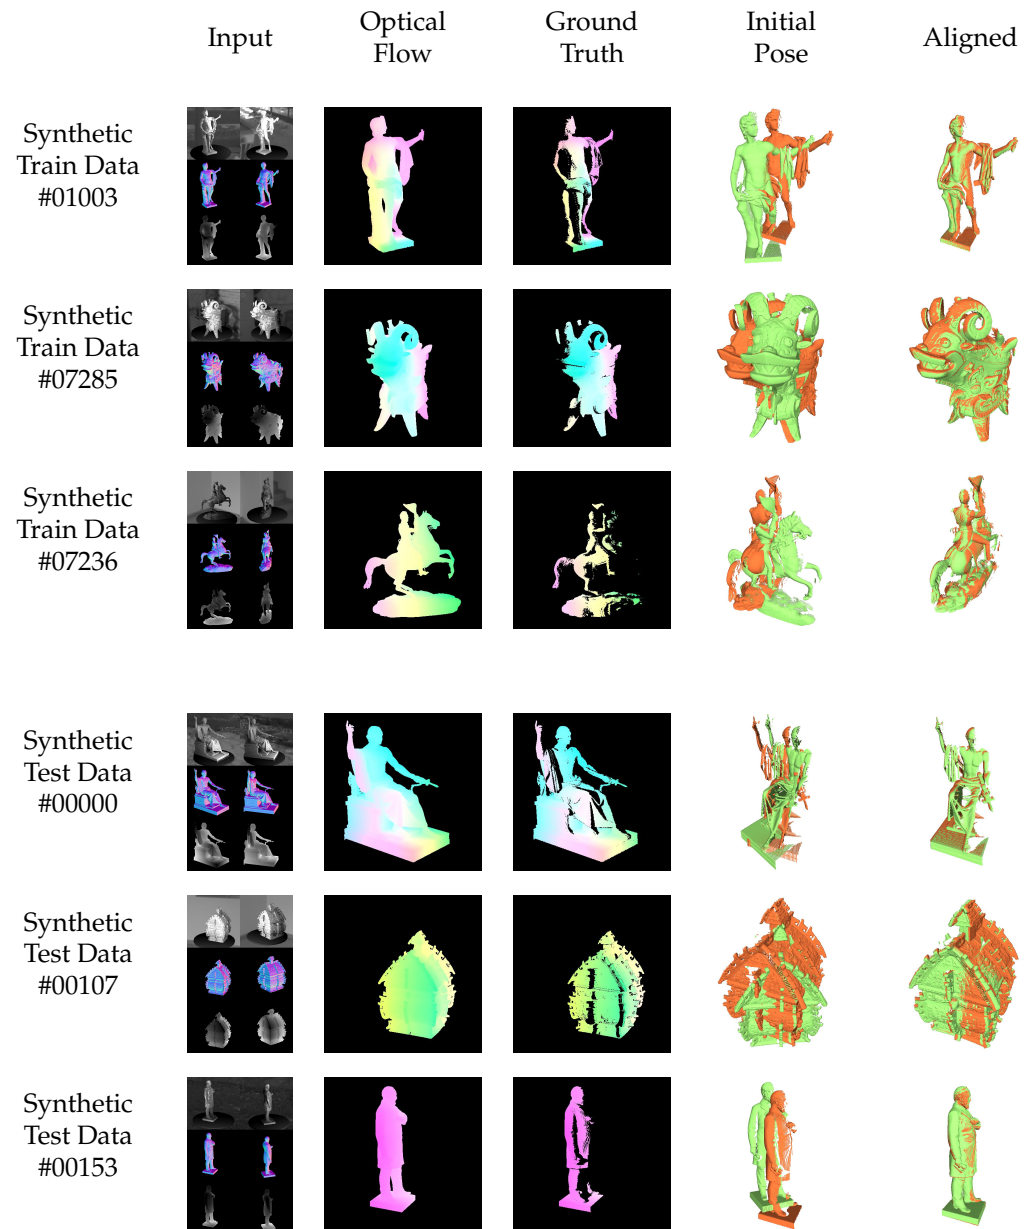


Figure 13. Qualitative results of the proposed method on training (top 3 rows) and test (bottom 3 rows) data of the synthetic *consistent light* dataset. The situation of *consistent light* represents the standard case, where, for example, the camera moves through a static scene with static light sources. The brightness assumption is usually not violated. The network generalizes well from known training to unknown test data.

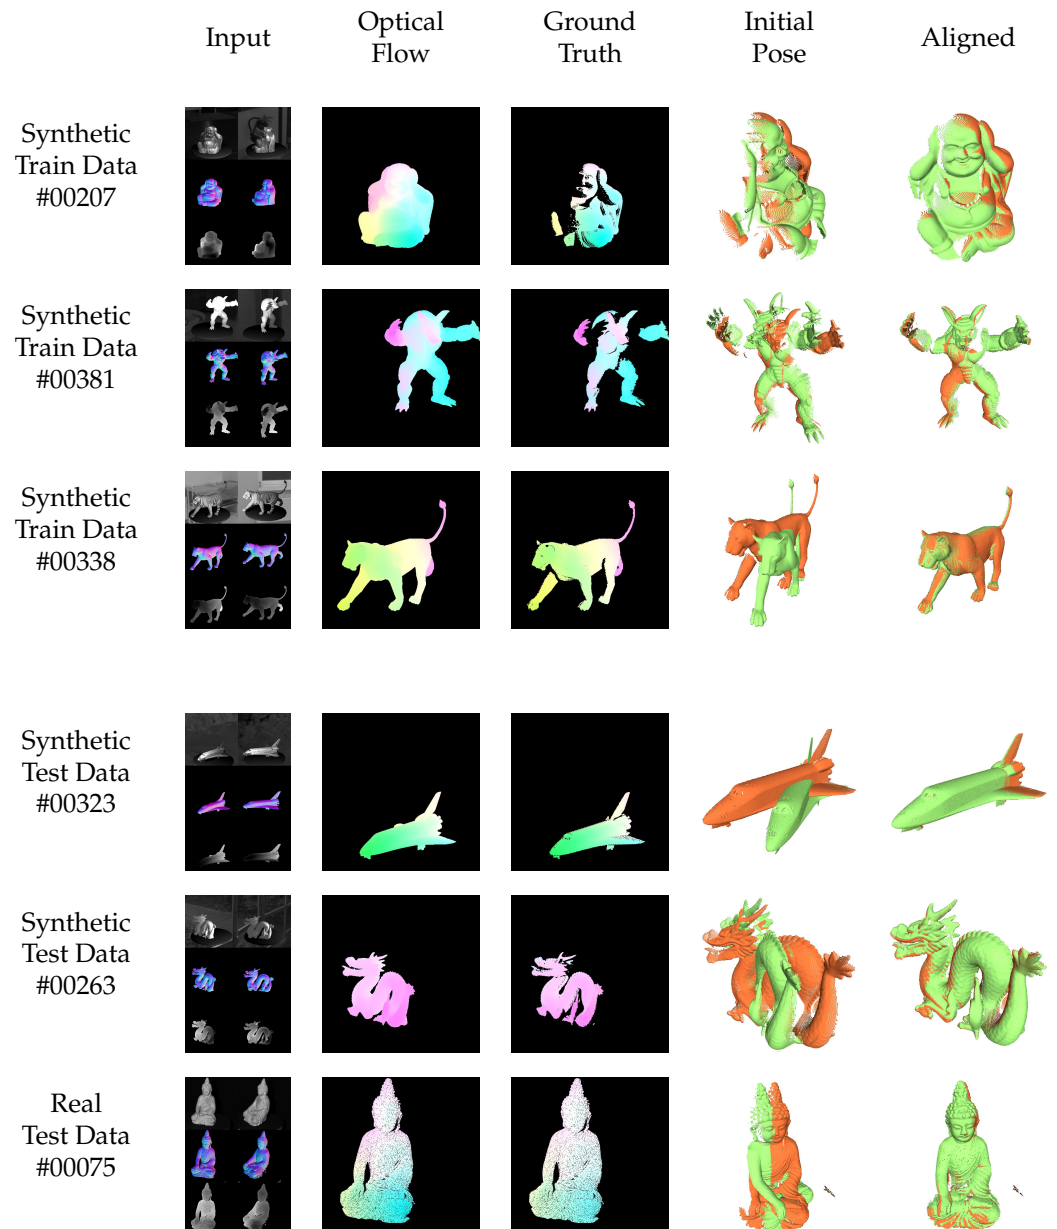


Figure 14. Qualitative results of the proposed method on training (top 3 rows) and test (bottom 3 rows) data of the synthetic *inconsistent light* dataset as well as real test data. The situation of *inconsistent light* represents the situation under investigation, motivating this paper, where the light sources or the objects in the scene move or rotate, yielding strong shading changes. The brightness assumption is dramatically violated. The network still generalizes well from known training to unknown test data. Even for real data without additional fine-tuning, the results are impressive.

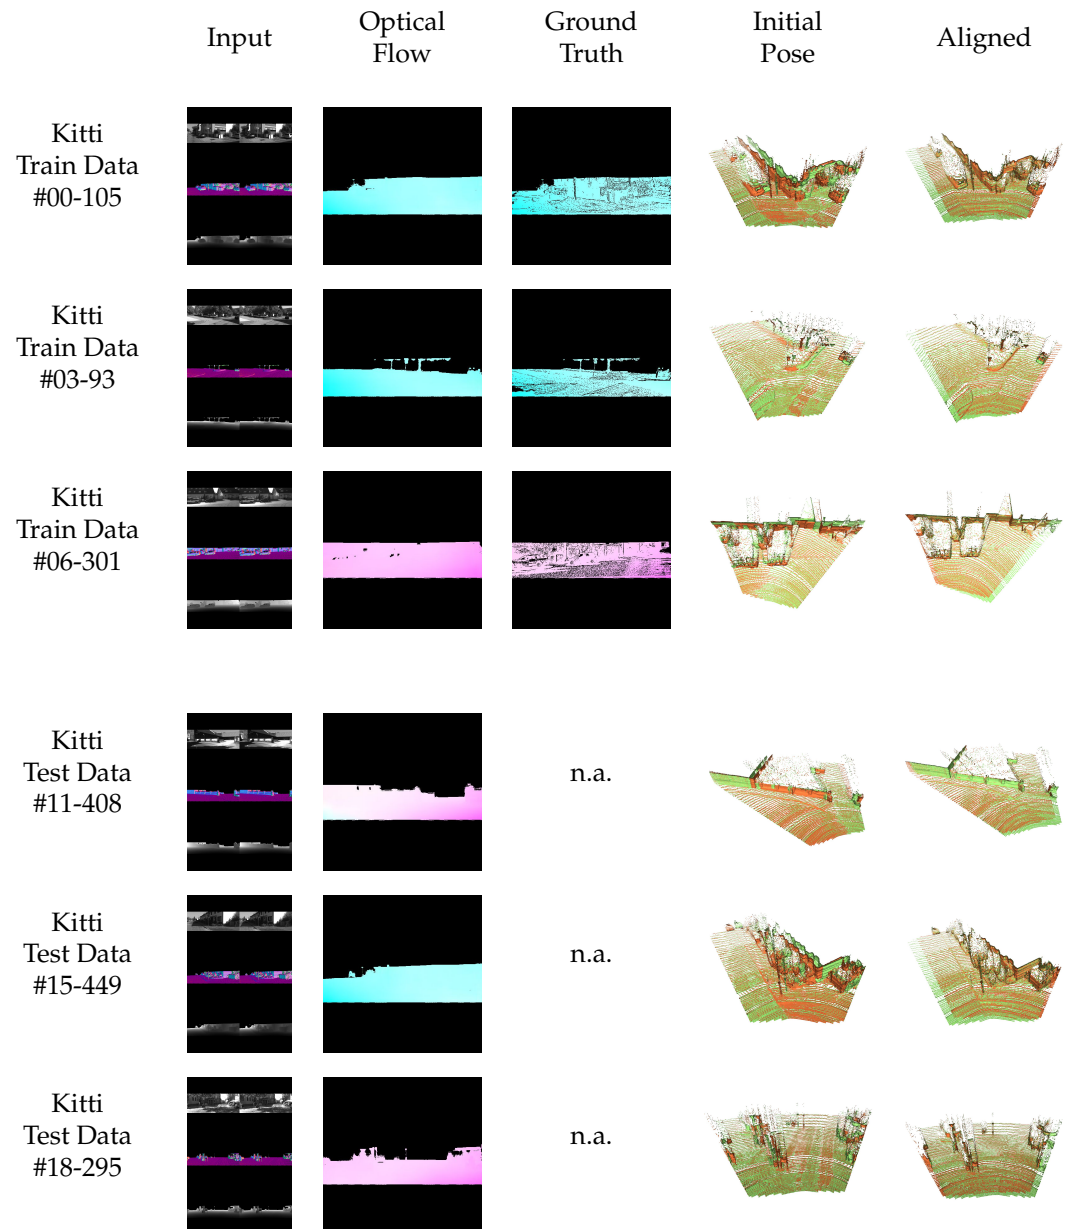


Figure 15. Qualitative results of the proposed method on training and test data of the *Kitti Odometry* dataset. The method also works on this kind of scenario with less rotations and less shading changes than in the mainly investigated case, but also handles noise resulting from the lidar depth measurement in the Kitti data. The network generalizes well from known training to unknown test data.

7.1. Quantitative Evaluation

For quantitative evaluation, we first compare the different architectures (*1 Step* and *3 Step*) on the datasets published together with this work. Table 1 shows the results on the full subsets with *consistent light* and *inconsistent light*. In both cases, the *3 Step* method yields superior results in comparison to the standard procedure that directly predicts rotation and translation jointly. In particular, the resulting rotation is much more accurate, resulting in an alignment error that is up to three times smaller than in the popular standard prediction method.

Table 1. Quantitative comparison of the 1 Step and the proposed 3 Step methods to predict the pose from given warped vertex and normal maps.

Light		Consistent		Inconsistent	
Data Type	Method	EPE	AE	EPE	AE
Train Data	1 Step	1.83	0.035	2.33	0.035
Train Data	3 Step	1.83	0.012	2.33	0.013
Test Data	1 Step	4.09	0.037	8.08	0.048
Test Data	3 Step	4.09	0.023	8.08	0.035

For completeness, we also trained the proposed architecture on the famous *Kitti Odometry* dataset. As mentioned, the data do not reflect the situations, where the strengths of the proposed architecture comes up. In addition, there are many procedures that are tuned to especially solve this common dataset. Nevertheless, our architecture is also able to deliver results that place within the ranking. Table 2 shows the methods that are also based on point clouds and therefore somehow comparable to the method presented. Our method would place around rank #100, which shows the opportunity of the method to be also applicable to other tasks.

Table 2. Extract of the ranking of the *Kitti Odometry* dataset showing point cloud based methods. The proposed method would be placed within the ranking, although rather at the end. Nevertheless, this shows the additional applicability of the method to other highly studied tasks.

Method	EPE	AE	R	t
V-LOAM (#2)	n.a.	n.a.	0.0013	0.54
LOAM (#3)	n.a.	n.a.	0.0013	0.55
GLIM (#5)	n.a.	n.a.	0.0015	0.59
CT-ICP (#6)	n.a.	n.a.	0.0014	0.59
SDV-LOAM (#7)	n.a.	n.a.	0.0015	0.60
CT-ICP2 (#8)	n.a.	n.a.	0.0012	0.60
wPICP (#11)	n.a.	n.a.	0.0015	0.62
FBLO (#12)	n.a.	n.a.	0.0014	0.62
HMLO (#14)	n.a.	n.a.	0.0014	0.62
filter-reg (#16)	n.a.	n.a.	0.0016	0.65
MULLS (#19)	n.a.	n.a.	0.0019	0.65
SMTD-LO (#22)	n.a.	n.a.	0.0020	0.66
PICP (#23)	n.a.	n.a.	0.0018	0.67
ELO (#24)	n.a.	n.a.	0.0021	0.68
IMLS-SLAM (#25)	n.a.	n.a.	0.0018	0.69
MC2SLAM (#26)	n.a.	n.a.	0.0016	0.69
ISC-LOAM (#28)	n.a.	n.a.	0.0022	0.72
Test-W (#30)	n.a.	n.a.	0.0033	0.79
PSF-LO (#31)	n.a.	n.a.	0.0032	0.82
S4-SLAM2 (#35)	n.a.	n.a.	0.0097	0.83
LIMO2_GP (#39)	n.a.	n.a.	0.0022	0.84
CAE-LO (#40)	n.a.	n.a.	0.0025	0.86

Table 2. Cont.

Method	EPE	AE	R	t
LIMO2 (#42)	n.a.	n.a.	0.0022	0.86
CPFG-slam (#44)	n.a.	n.a.	0.0025	0.87
SD-DEVO (#49)	n.a.	n.a.	0.0028	0.88
PNDT LO (#50)	n.a.	n.a.	0.0030	0.89
LIMO (#58)	n.a.	n.a.	0.0026	0.93
SuMa-MOS (#67)	n.a.	n.a.	0.0033	0.99
SuMa++ (#69)	n.a.	n.a.	0.0034	1.06
DEMO (#74)	n.a.	n.a.	0.0049	1.14
STEAM-L WNOJ (#83)	n.a.	n.a.	0.0058	1.22
LiViOdo (#84)	n.a.	n.a.	0.0042	1.22
STEAM-L (#87)	n.a.	n.a.	0.0061	1.26
SALO (#93)	n.a.	n.a.	0.0051	1.37
SuMa (#95)	n.a.	n.a.	0.0034	1.39
Flow2PoseNet 3 Step	1.18	0.019	0.0019	2.73
Deep-CLR (#134)	n.a.	n.a.	0.0104	3.83
SLL (#163)	n.a.	n.a.	0.2645	90.05

7.2. Predicted Dense Optical Flow


A special feature of the proposed method is its coarse to fine pyramidal optical flow base, combined with the rigid pose extraction. Therefore, one can assume that the optical flow predicting sub-network learns rigidity relations from the extractability of the rigid pose from the dense optical flow. As shown in Figure 16, the ground truth optical flow (column 2) that has been used for training and evaluating the networks is sparse, as it only contains the flow of points that are visible in both views. As the data are created synthetically, it is possible to also render dense ground truth optical flows (column 4) that contain the flow of points that are occluded in one of the views and therefore may not be computable at all by the network. As can be seen, the predicted optical flow (column 3) is dense. It also predicts flow values for points that are not visible in both views. These values result from the context of other points, where the flow can be estimated stably. The network learns how the flow behaves for rigid objects and transfers the knowledge to interpolated pixels. This works as well for objects that are known from the training set (rows 1 and 3), test objects that have never been used for training (rows 2 and 4), the *ConsistentLight* case (rows 1 and 2) as well as for the *InconsistentLight* case (rows 3 and 4). Table 3 moreover shows that the resulting *Endpoint Errors (EPE)* do not dramatically increase for the invisible points, which indicates that the network learns to predict flows for the invisible points from context, according to the behavior of rigid objects.

Table 3. Quantitative results for the visible and invisible points in the evaluated scenes. The resulting *Endpoint Errors (EPE)* do not heavily increase. The network is still able to predict accurate flows from the context of visible points and to generalize to the test data as well as for the consistent and inconsistent data.


Consistent Light		
Data Type	Visible Points EPE	Invisible Points EPE
Train Data	2.7446	3.4978
Test Data	3.6411	4.9284
Inconsistent Light		
Data Type	Visible Points EPE	Invisible Points EPE
Train Data	3.6974	5.4024
Test Data	4.7996	4.7703




(a) Consistent Light: Train Object



(b) Consistent Light: Test Object



(c) Inconsistent Light: Train Object



(d) Inconsistent Light: Test Object

Figure 16. Qualitative results of the predicted (dense) optical flow. The network allows for computing accurate flows for invisible pixels from the context of visible parts as well as for the consistent and inconsistent data.

8. Conclusions

In this paper, a method has been presented that combines optical flow estimation of rigid scenes with a posterior pose estimation. In this way, including several contributions, a method has been developed that allows scenes with difficult lighting conditions to be registered in a stable way.

Optical flow is thereby estimated accurately using geometric, shading and texture features. The variety of different feature types allows the system to be trained to be illumination resistant (using geometric and normal features) without having to completely sacrifice potentially important texture features.

The pose is then stably estimated from the warped normals and vertex maps using a new 3-step procedure. This has, compared to typical approaches that directly infer the pose, significant advantages especially in cases with strong rotations that often cause the considered shading changes.

The combination of optical flow and rigid pose estimation allows the pose to benefit from the features of different levels of the underlying coarse-to-fine flow approach, which means that the method is not dependent on highly accurate features and can also align smooth scenes with weak features. In turn, the optical flow sub-network learns a typical flow behavior of rigid scenes from the posterior estimability of the pose. This allows accurate dense estimates to be achieved, even for occluded areas based on context and overall learned behavior.

Author Contributions: Conceptualization, T.F.; methodology, T.F.; software, T.F.; validation, T.F.; formal analysis, T.F.; investigation, T.F.; resources, T.F.; data curation, T.F.; writing—original draft preparation, T.F.; writing—review and editing, T.F. and G.R.; visualization, T.F.; supervision, G.R.; project administration, G.R.; funding acquisition, G.R. and D.S. All authors have read and agreed to the published version of the manuscript.

Funding: This work was partially funded by the Federal Ministry of Education and Research Germany under the project DECODE (01IW21001).

Conflicts of Interest: The authors declare no conflict of interest.

References

1. Ferraz, L.; Binefa, X.; Moreno-Noguer, F. Very fast solution to the PnP problem with algebraic outlier rejection. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, Columbus, OH, USA, 23–28 June 2014; pp. 501–508.
2. Horn, B.K.; Schunck, B.G. Determining optical flow. *Artif. Intell.* **1981**, *17*, 185–203.
3. Čech, J.; Sanchez-Riera, J.; Horaud, R. Scene flow estimation by growing correspondence seeds. In Proceedings of the CVPR 2011, IEEE, Washington, DC, USA, 20–25 June 2011; pp. 3129–3136.
4. Huguet, F.; Devernay, F. A variational method for scene flow estimation from stereo sequences. In Proceedings of the 2007 IEEE 11th International Conference on Computer Vision, IEEE, Daejeon, Korea, 5–6 November 2007; pp. 1–7.
5. Isard, M.; McCormick, J. Dense motion and disparity estimation via loopy belief propagation. In Proceedings of the Asian Conference on Computer Vision, Springer, Hyderabad, India, 13–16 January 2006; pp. 32–41.
6. Li, R.; Sclaroff, S. Multi-scale 3D scene flow from binocular stereo sequences. *Comput. Vis. Image Underst.* **2008**, *110*, 75–90.
7. Basha, T.; Moses, Y.; Kiryati, N. Multi-view scene flow estimation: A view centered variational approach. *Int. J. Comput. Vis.* **2013**, *101*, 6–21.
8. Park, J.; Oh, T.H.; Jung, J.; Tai, Y.W.; Kweon, I.S. A tensor voting approach for multi-view 3D scene flow estimation and refinement. In Proceedings of the European Conference on Computer Vision, Springer, Florence, Italy, 7–13 October 2012; pp. 288–302.
9. Zhang, X.; Chen, D.; Yuan, Z.; Zheng, N. Dense scene flow based on depth and multi-channel bilateral filter. In Proceedings of the Asian Conference on Computer Vision, Springer, Daejeon, Korea, 5–9 November 2012; pp. 140–151.
10. Ferstl, D.; Reinbacher, C.; Riegler, G.; Rütger, M.; Bischof, H. aTGV-SF: Dense variational scene flow through projective warping and higher order regularization. In Proceedings of the 2014 2nd International Conference on 3D Vision, IEEE, Tokyo, Japan, 8–11 December 2014; Volume 1, pp. 285–292.
11. Letouzey, A.; Petit, B.; Boyer, E. Scene flow from depth and color images. In Proceedings of the BMVC 2011-British Machine Vision Conference, BMVA Press, Dundee, Scotland, 30 August–1 September 2011, pp. 46–1.
12. Gottfried, J.M.; Fehr, J.; Garbe, C.S. Computing range flow from multi-modal kinect data. In Proceedings of the International Symposium on Visual Computing, Springer, Las Vegas, NV, USA, 26–28 September 2011; pp. 758–767.
13. Herbst, E.; Ren, X.; Fox, D. Rgb-d flow: Dense 3-d motion estimation using color and depth. In Proceedings of the 2013 IEEE International Conference on Robotics and Automation, IEEE, Kagawa, Japan, 6–10 May 2013; pp. 2276–2282.
14. Quiroga, J.; Brox, T.; Devernay, F.; Crowley, J. Dense semi-rigid scene flow estimation from rgbd images. In Proceedings of the European Conference on Computer Vision, Springer, Zurich, Switzerland, 5–12 September 2014; pp. 567–582.
15. Dosovitskiy, A.; Fischer, P.; Ilg, E.; Hausser, P.; Hazirbas, C.; Golkov, V.; Van Der Smagt, P.; Cremers, D.; Brox, T. FlowNet: Learning optical flow with convolutional networks. In Proceedings of the IEEE international Conference on Computer Vision, Santiago, Chile, 7–13 December 2015; pp. 2758–2766.

16. Ilg, E.; Mayer, N.; Saikia, T.; Keuper, M.; Dosovitskiy, A.; Brox, T. FlowNet 2.0: Evolution of optical flow estimation with deep networks. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, Honolulu, HI, USA, 21–26 July 2017; pp. 2462–2470.
17. Sun, D.; Yang, X.; Liu, M.Y.; Kautz, J. Pwc-net: Cnns for optical flow using pyramid, warping, and cost volume. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, Salt Lake City, UT, USA, 18–23 June 2018; pp. 8934–8943.
18. Hur, J.; Roth, S. MirrorFlow: Exploiting symmetries in joint optical flow and occlusion estimation. In Proceedings of the IEEE International Conference on Computer Vision, Venice, Italy, 22–29 October 2017; pp. 312–321.
19. Meister, S.; Hur, J.; Roth, S. Unflow: Unsupervised learning of optical flow with a bidirectional census loss. In Proceedings of the AAAI Conference on Artificial Intelligence, New Orleans, LA, USA, 2–7 February 2018; Volume 32.
20. Wang, Y.; Yang, Y.; Yang, Z.; Zhao, L.; Wang, P.; Xu, W. Occlusion aware unsupervised learning of optical flow. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, Salt Lake City, UT, USA, 18–23 June 2018; pp. 4884–4893.
21. Hur, J.; Roth, S. Iterative residual refinement for joint optical flow and occlusion estimation. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, Long Beach, CA, USA, 15–20 June 2019; pp. 5754–5763.
22. Hui, T.W.; Tang, X.; Loy, C.C. LiteflowNet: A lightweight convolutional neural network for optical flow estimation. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, Salt Lake City, UT, USA, 18–23 June 2018; pp. 8981–8989.
23. Hui, T.W.; Loy, C.C. LiteflowNet3: Resolving correspondence ambiguity for more accurate optical flow estimation. In Proceedings of the European Conference on Computer Vision, Springer, Glasgow, UK, 23–28 August 2020; pp. 169–184.
24. Liu, P.; Lyu, M.; King, I.; Xu, J. SelfFlow: Self-supervised learning of optical flow. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, Long Beach, CA, USA, 15–20 June 2019; pp. 4571–4580.
25. Jonschkowski, R.; Stone, A.; Barron, J.T.; Gordon, A.; Konolige, K.; Angelova, A. What matters in unsupervised optical flow. In Proceedings of the European Conference on Computer Vision, Springer, Glasgow, UK, 23–28 August 2020; pp. 557–572.
26. Zhu, A.Z.; Yuan, L.; Chaney, K.; Daniilidis, K. Unsupervised event-based learning of optical flow, depth, and egomotion. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, Long Beach, CA, USA, 15–20 June 2019; pp. 989–997.
27. Janai, J.; Guney, F.; Ranjan, A.; Black, M.; Geiger, A. Unsupervised learning of multi-frame optical flow with occlusions. In Proceedings of the European Conference on Computer Vision (ECCV), Munich, Germany, 8–14 September 2018; pp. 690–706.
28. Tu, Z.; Xie, W.; Zhang, D.; Poppe, R.; Veltkamp, R.C.; Li, B.; Yuan, J. A survey of variational and CNN-based optical flow techniques. *Signal Process. Image Commun.* **2019**, *72*, 9–24.
29. Rishav, R.; Battrawy, R.; Schuster, R.; Wasenmüller, O.; Stricker, D. DeepLiDARFlow: A Deep Learning Architecture For Scene Flow Estimation Using Monocular Camera and Sparse LiDAR. In Proceedings of the 2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), IEEE, Las Vegas, NV, USA, 25–29 October 2020; pp. 10460–10467.
30. Eldesokey, A.; Felsberg, M.; Khan, F.S. Confidence propagation through cnns for guided sparse depth regression. *IEEE Trans. Pattern Anal. Mach. Intell.* **2019**, *42*, 2423–2436.
31. Yang, Z.; Wang, P.; Wang, Y.; Xu, W.; Nevatia, R. Every pixel counts: Unsupervised geometry learning with holistic 3d motion understanding. In Proceedings of the European Conference on Computer Vision (ECCV) Workshops, Munich, Germany, 8–14 October 2018.
32. Yin, Z.; Shi, J. Geonet: Unsupervised learning of dense depth, optical flow and camera pose. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, Salt Lake City, UT, USA, 18–23 June 2018; pp. 1983–1992.
33. Zou, Y.; Luo, Z.; Huang, J.B. Df-net: Unsupervised joint learning of depth and flow using cross-task consistency. In Proceedings of the European Conference on Computer Vision (ECCV), Munich, Germany, 8–14 September 2018; pp. 36–53.
34. Kendall, A.; Grimes, M.; Cipolla, R. PoseNet: A convolutional network for real-time 6-dof camera relocalization. In Proceedings of the IEEE International Conference on Computer Vision, Santiago, Chile, 7–13 December 2015; pp. 2938–2946.
35. Vijayanarasimhan, S.; Ricco, S.; Schmid, C.; Sukthankar, R.; Fragkiadaki, K. Sfm-net: Learning of structure and motion from video. *arXiv* **2017**, arXiv:1704.07804.
36. Zhou, T.; Brown, M.; Snavely, N.; Lowe, D.G. Unsupervised learning of depth and ego-motion from video. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, Honolulu, HI, USA, 21–26 July 2017; pp. 1851–1858.
37. Zhang, Z.; Dai, Y.; Sun, J. Deep learning based point cloud registration: an overview. *Virtual Real. Intell. Hardw.* **2020**, *2*, 222–246.
38. Villena-Martinez, V.; Oprea, S.; Saval-Calvo, M.; Azorin-Lopez, J.; Fuster-Guillo, A.; Fisher, R.B. When deep learning meets data alignment: A review on deep registration networks (drns). *Appl. Sci.* **2020**, *10*, 7524.
39. Fragkiadaki, K.; Hu, H.; Shi, J. Pose from flow and flow from pose. In Proceedings of the Conference on Computer Vision and Pattern Recognition, Portland, OR, USA, 23–28 June 2013; pp. 2059–2066.
40. Piga, N.A.; Onyshchuk, Y.; Pasquale, G.; Pattacini, U.; Natale, L. ROFT: Real-Time Optical Flow-Aided 6D Object Pose and Velocity Tracking. *IEEE Robot. Autom. Lett.* **2021**, *7*, 159–166.
41. Holzer, S.; Rusu, R.B.; Dixon, M.; Gedikli, S.; Navab, N. Adaptive neighborhood selection for real-time surface normal estimation from organized point cloud data using integral images. In Proceedings of the 2012 IEEE/RSJ International Conference on Intelligent Robots and Systems, IEEE, Algarve, Portugal, 7–12 October 2012; pp. 2684–2689.
42. Eldesokey, A.; Felsberg, M.; Khan, F.S. Propagating confidences through cnns for sparse data regression. *arXiv* **2018**, arXiv:1805.11913.
43. Unity Game Engine. Available online: <https://unity.com> (accessed on 2 December 2021).

-
44. Stanford Scanning Repository. Available online: <http://graphics.stanford.edu/data/3Dscanrep/> (accessed on 10 March 2022).
 45. Zhou, K.; Wang, X.; Tong, Y.; Desbrun, M.; Guo, B.; Shum, H.Y. Texturemontage. In *ACM SIGGRAPH 2005 Papers*; Association for Computing Machinery: New York, NY, USA, 2005; pp. 1148–1155.
 46. Smithsonian 3D Digitization. Available online: <https://3d.si.edu/> (accessed on 10 March 2022).