



HAL
open science

The Right Spin: Learning Object Motion from Rotation-Compensated Flow Fields

Pia Bideau, Erik Learned-Miller, Cordelia Schmid, Karteek Alahari

► **To cite this version:**

Pia Bideau, Erik Learned-Miller, Cordelia Schmid, Karteek Alahari. The Right Spin: Learning Object Motion from Rotation-Compensated Flow Fields. 2022. hal-03593853

HAL Id: hal-03593853

<https://inria.hal.science/hal-03593853v1>

Preprint submitted on 2 Mar 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

The Right Spin: Learning Object Motion from Rotation-Compensated Flow Fields

Pia Bideau^{1*}, Erik Learned-Miller², Cordelia Schmid³ and Karteek Alahari⁴

^{1*}Science of Intelligence, TU Berlin, Marchstr, Berlin, 10587, Germany.

²University of Massachusetts, Governors Dr, Amherst, 01002, MA, US.

³Inria, École Normale Supérieure, CNRS, PSL Research University, Paris, France.

⁴Univ. Grenoble Alpes, Inria, CNRS, Grenoble INP, LJK, 38000, Grenoble, France.

*Corresponding author(s). E-mail(s): p.bideau@tu-berlin.de;

Abstract

Both a good understanding of geometrical concepts and a broad familiarity with objects lead to our excellent perception of moving objects. The human ability to detect and segment moving objects works in the presence of multiple objects, complex background geometry, motion of the observer and even camouflage. How humans perceive moving objects so reliably is a longstanding research question in computer vision and borrows findings from related areas such as psychology, cognitive science and physics. One approach to the problem is to teach a deep network to model all of these effects. This contrasts with the strategy used by human vision, where cognitive processes and body design are tightly coupled and each is responsible for certain aspects of correctly identifying moving objects. Similarly from the computer vision perspective, there is evidence (Bideau et al., 2018; Irani and Anandan, 1998) that classical, geometry-based techniques are better suited to the “motion-based” parts of the problem, while deep networks are more suitable for modeling appearance. In this work, we argue that the coupling of camera rotation and camera translation can create complex motion fields that are difficult for a deep network to untangle directly. We present a novel probabilistic model to estimate the camera’s rotation given the motion field. We then rectify the flow field to obtain a rotation-compensated motion field for subsequent segmentation. This strategy of first estimating camera motion, and then allowing a network to learn the remaining parts of the problem, yields improved results on the widely used DAVIS benchmark as well as the recently published motion segmentation data set MoCA (Moving Camouflaged Animals).

Keywords: Motion segmentation, Video segmentation, Optical flow, Camera motion estimation

1 Introduction

The human visual system has the ability to detect independently moving objects within a high variety of different environments. While we are moving through the world our eye captures a large amount of visual information over time. Often, we are not aware of the remarkable preprocessing steps that happen almost unnoticed. For example, human eye movements induce two major simplifications to incoming

images before visual information is processed by the visual cortex. These are (1) stabilizing the image - reducing the amount of local change due to motion, and (2) changing the direction of gaze (Walls, 1962; Longuet-Higgins et al., 1980).

Here, we revisit this approach to motion segmentation that separates the problem into two parts: first, we preprocess the perceived motion field following well known geometrical concepts leading to important simplifications and second, learning to segment

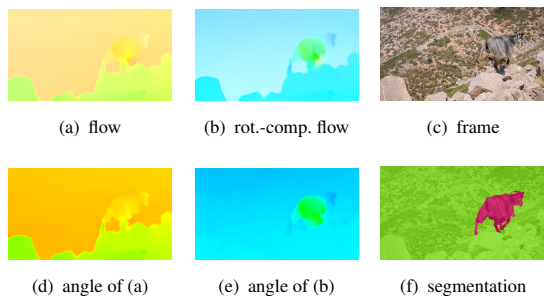


Fig. 1 What is moving? Coupling of camera rotation and camera translation often create complex motion fields that are difficult for a network to untangle. Instead we propose a strategy to learn object motion patterns based on rotation compensated flow.

independently moving objects from these simplified motion fields.

In computer vision, the task of motion segmentation attempts to analyze the perceived motion and to segment a video sequence into static environment (if any) and independently moving objects (Bideau and Learned-Miller, 2016a). Interpreting the motion field accurately, and then drawing the right conclusions about what is moving in the world and what is static, is a complex process. Even in biological vision systems applied strategies are still only partially understood.

Unlike most end-to-end learning-based approaches, where a model learns all necessary steps between the input and the final output, we break down the problem of motion segmentation into two sub-problems: adjusting the optical flow to remove the effects of camera rotation (*rotation compensation*) using classical approaches based on perspective projection and learning to segment the remaining optical flow into static background and moving objects. The step of compensating for camera rotation is a challenging one, since the flow field is only a noisy estimate of the motion field (Bideau et al., 2018; Bideau and Learned-Miller, 2016b). In cases of little motion or of featureless areas, the observed flow field is often erroneous and thus the true camera motion and object motion is hard to estimate accurately.

To this end, we present a novel probabilistic method for estimating camera rotation and derive a new likelihood function modeling the probability of an observed optical flow field, given our estimated (ideal) motion field. A CNN framework is then integrated for learning to segment moving objects after the motion of the camera has been determined.

Our contributions include: (i) estimating the camera rotation and translational motion direction in the

presence of moving objects, using a new likelihood maximization approach, (ii) given the rotation compensated flow, we show that the task of learning motion patterns is improved, resulting in a better motion segmentation performance shown on two data sets - the widely used DAVIS benchmark (Perazzi et al., 2016) and the recently published data set MoCA (Lamdouar et al., 2020).

The paper is structured as follows. In Section 2 we review relevant work on motion segmentation starting from classical geometry based approaches and concluding with the most recent work using convolutional neural networks to segment moving objects from optical flow. In Section 3, we develop an end-to-end approach for motion segmentation. We briefly review the basics about the motion field and how it is related to camera motion, depth and object motion (Section 3.1). Building upon key concepts of perspective projection the methodological approach is derived in two subsections: estimating the camera rotation to produce rotation-compensated flow fields (Section 3.2) and segmenting the remaining (noisy) translational flow field into independently moving objects and static background (Section 3.3). A multi-faceted evaluation of the proposed approach, including multiple ablation studies has been carried out and is shown in Experiments (Section 4).

2 Related Work

Many works tackling the problem of motion segmentation focus on *binary motion segmentation*, where pixels are classified as either moving or being part of the static background. In that case no distinction is made between differently moving objects (Bideau and Learned-Miller, 2016b; Narayana et al., 2013; Papazoglou and Ferrari, 2013; Faktor and Irani, 2014). Others (Taylor et al., 2015; Keuper et al., 2015; Fragkiadaki et al., 2012) address *multi-label motion segmentation*, where a separate label is given to each independently moving object. Our work tackles binary motion segmentation, but we consider both views onto the segmentation problem in this review of related work.

2.1 Classical approaches

Methods based on feature clustering

To capture motion information, typically point trajectories are either formed by tracked image features

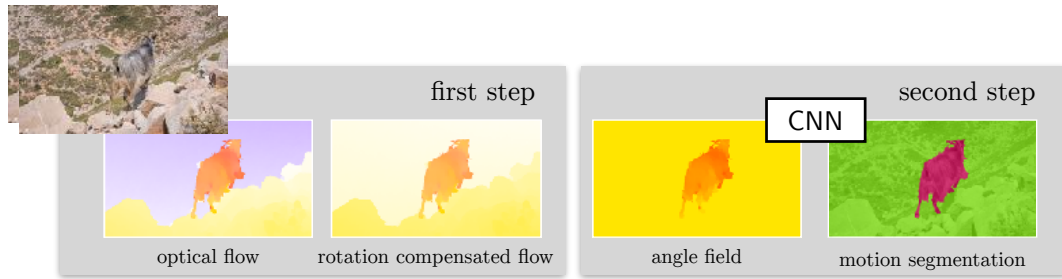


Fig. 2 Getting the right spin. We first compensate the observed motion field for camera rotation (“first step”), and segment the remaining translational optical flow field using a learning based approach (“second step”). The observed flow field on the left has complex motion patterns: the motion directions of foreground and background are pointing in opposite directions, due to large variance in scene depth, and the combined impact of camera rotation and translation. Estimating the camera rotation (“the right spin”), and compensating the flow field for this rotation simplifies the motion field dramatically, in this case yielding similar motion directions for foreground and background. This provides simpler inputs to our learning based motion segmentation framework.

or dense optical flow. Then trajectories sharing similar motion characteristics are grouped into coherent motion clusters describing the motion of a particular object (Keuper et al., 2015; Brox and Malik, 2010; Fragkiadaki et al., 2012; Ochs and Brox, 2011; Keuper, 2017; Yan and Pollefeys, 2006; Shen et al., 2018; Lezama et al., 2011).

These approaches vary in defining typical motion characteristics for clustering. Yan et. al (Yan and Pollefeys, 2006) propose to cluster trajectories based on geometric constraints (trajectories of the same motion lie in a manifold) and locality. In (Keuper, 2017; Keuper et al., 2015) the segmentation problem is represented as a minimum cost multicut graph problem, where edge weights are computed from motion, position and color cues.

These trajectory based clustering approaches reach their limit if understanding of the scene structure is necessary to segment a moving object correctly. Trajectories perfectly represent long-term pixel displacements between a sequence of frames. Pixel displacements however are a function of depth and motion. Thus trajectory based clustering methods often form clusters not only for independently moving objects, but also for objects at different depths. For instance if the camera is translating and rotating rocks close to the camera produce a very different flow pattern that the far away scene (see Figure 2), thus those two areas would form two separate clusters although neither the rock nor the far away scene is moving.

Methods based on occlusions (Ogale et al., 2005; Taylor et al., 2015) are subject to similar depth-related problems, since occlusions could be caused at depth boundaries as well as motion boundaries. A distinction is often not made.

Methods based on projective geometry

Projective geometry is an extension of the Euclidean and affine space and contains properties of perspective projection. It is widely used as a mathematical formalism to describe the geometry of cameras and its associated transformations (Torr, 1998; Zamalieva and Yilmaz, 2014; Wang and Adelson, 1994; Ke and Kanade, 2002; Jin et al., 2008; Xiao and Shah, 2005; Vidal and Ma, 2004; Xu et al., 2018).

Different from trajectory based clustering methods, motion segmentation approaches relying on projective geometry analyze the optical flow between a pair of frames, grouping pixels into regions where flow is consistent with motion models that are explainable by projective geometry (Torr, 1998; Zamalieva and Yilmaz, 2014; Wang and Adelson, 1994; Ke and Kanade, 2002; Jin et al., 2008; Xiao and Shah, 2005; Xu et al., 2018). Torr (Torr, 1998) develops a sophisticated probabilistic model of optical flow, building a mixture model that explains an arbitrary number of rigid components within the scene. Interestingly, he assigns different types of motion models to each object based on model fitting criteria. Zamalieva et al. (Zamalieva and Yilmaz, 2014) and Xun Xu et al. (Xu et al., 2018) present a combination of methods that rely on both - projective geometry (homography estimation) and perspective projection (fundamental matrix estimation). The two methods have complementary strengths, and the authors attempt to select among the best dynamically.

Methods relying on projective geometry perform well in cases of planar motion (motion obtained by a translating or rotating camera picturing a planar scene or a very distant scene, where effects of 3D parallax are negligible), however similarly to cluster

based approaches these methods fall short in case of complex scene geometry.

Horn identified specific drawbacks of using projective geometry in such estimation problems and has argued that methods based directly on perspective projection are less prone to overfitting in the presence of noise (Horn, 1999).

Methods based on perspective projection

Perspective geometry allows us to mathematically explain and model the process of how the three-dimensional world is projected on to a just two-dimensional image plane. Artists and scientists like Alberti, Brunelleschi, Dürer and da Vinci studied effects of perspective projection about 500 years back in time (Pirenne, 1952). These insights have made a significant contribution to current successes in computer vision. One of the key aspects of perspective projection is the observation that two parallel lines (in the euclidean space) are transformed to two lines that intersect in the vanishing point at the horizon on the image plane.

It has been shown that motion segmentation approaches based on perspective projection (Irani and Anandan, 1998; Bideau and Learned-Miller, 2016b; Bideau et al., 2018; Narayana et al., 2013; Vidal et al., 2002; Zhang et al., 2007; Yang and Ramanan, 2021) are more accurate (in terms of model agreement to the physical world) than those based on projective geometry, since the latter omits certain constraints in modeling image transformations (Horn, 1999; Bideau and Learned-Miller, 2016b). Having a model that is confirm with the physical world might be especially critical for tasks where interaction with the physical world is required in a second step such as in robotics or autonomous driving.

2.2 Learning motion segmentation using convolutional neural networks

Methods based on supervised learning

Several approaches as (Tokmakov et al., 2017a,b; Jain et al., 2017; Cheng et al., 2017a; Dave et al., 2019; Ranjan et al., 2019; Vertens et al., 2017; Mahadevan et al., 2020; Lamdouar et al., 2020; Cheng et al., 2017b) have explored the strength of deep neural networks to learn motion patterns of moving objects and to produce binary motion masks distinguishing whether a pixel belongs to a moving object or not. Most approaches propose a two-stream architecture (Tokmakov et al., 2017b; Jain et al., 2017;

Dave et al., 2019) to separately process motion and appearance.

These approaches learn motion patterns given the optical flow, the raw video frames or optical flow together video frames. Rather than following the true physics of image formation, convolutional neural networks are able to learn high level motion patterns of background motion and object motion. This ability has the clear advantage of not being dependent upon technical camera parameters such as the focal length or image distortions due to various lens characteristics or constraints induced by technical parts of the camera (mechanical or electronic).

Methods based on self-supervised learning

General concerns of deep-learning based approaches and in particular supervised approaches are overfitting to a particular type of object category that is likely to move (Dave et al., 2019) and the lack of large amounts of training data. To overcome the problem of limited training data, two straight forward approaches are either using synthetic training data (Tokmakov et al., 2017a,b) or relying on noisy estimates of the motion field (Jain et al., 2017) using other algorithms (Sun et al., 2018; Ilg et al., 2017; Sun et al., 2010). However, both paths are still in need of large amounts of training data (although no additional manual annotations are required in these cases), this rises the need for self-supervised approaches (Yang et al., 2021; Lu et al., 2019a; Yang et al., 2019; Lai et al., 2020; Gordon et al., 2019; Bideau et al., 2018). Incorporating knowledge about the real world physics into the training procedure of a neural network is an alternative to various kinds of data augmentation approaches that is subject of current research (Tung et al., 2019; Yang and Ramanan, 2021). Some of those ideas have been already successfully applied in context of self-supervised learning (Gordon et al., 2019; Bideau et al., 2018).

Here, we propose a novel approach to the motion segmentation problem that specifically combines aspects of perspective projection and learns general object motion patterns.

3 Learning object motion from rotation-compensated flow

As most of the previous works we define a *moving object* as a *collection of matter that independently moves as a whole in the 3D world*. An overview of

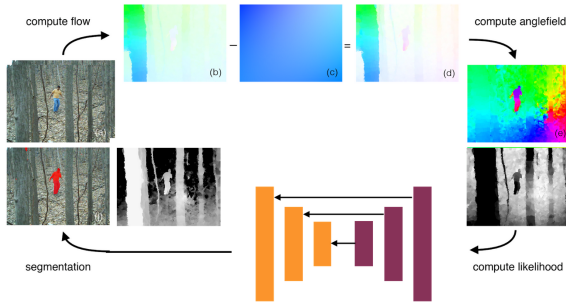


Fig. 3 Overview of our approach. Given the optical flow (b) the camera rotation is estimated (Section 3.2.2). The flow \vec{v}_r due to camera rotation is defined by the motion parameters (A, B, C) . (c) is subtracted from the optical flow $\vec{\sigma}$ to produce a translational flow $\vec{\sigma}_t$. The flow angle $\theta_{\vec{\sigma}_t}$ and magnitudes $|\vec{\sigma}_t|$ are shown in (e).

our approach for motion segmentation is shown in Figure 3. Given an estimate of the motion field (optical flow) each frame is segmented into static environment and independently moving objects. To achieve this we present an approach where we first estimate the camera rotation and then use this knowledge to form a rotation-compensated flow field. A network is trained that takes rotation-compensated flow fields as input and outputs motion segmentation masks. To this end, we combine our novel geometry-based method for estimating camera rotation, and a CNN framework for learning to segment moving objects.

In the following we will revise relevant background information about the formation of a motion field, that occurs on the camera sensor as the camera moves (Section 3.1). Building on this, we propose a novel approach to estimate camera rotation in complex environments, considering scene depth as well as independently moving objects (Section 3.2). In Section 3.3, we propose an approach similar to (Bideau et al., 2018) that learns to segment the rotation compensated motion field into static background and independently moving objects.

3.1 The Motion Field: A Geometrical Analysis

The motion field captures pixel displacements between two consecutive frames. Displacements arise typically due to one of the following factors: (1) a moving camera, (2) one or more objects moving in the 3D world. These pixel displacements depend not only on the speed of objects or the camera, but also the scene geometry.

As an example to illustrate the different factors that influence the formation of the motion field, let’s

consider the “goat” sequence from the DAVIS data set (Figure 2). Based on the original flow field it is hard to estimate which pixels belong to the moving object and which belong to static background. The direction of the flow in the background region differs significantly from the flow describing the motion of the rocks in the foreground region (motion direction is color encoded). However, neither the background nor the rocks are moving differently in the 3D world. To detect objects that are actually moving independently in 3D it is necessary to decompose the observed motion field. We formalize these observations and review the geometrical construction of the motion field.

3.1.1 Motion field

Let $[U, V, W]$ be the parameters describing the camera translation and $[A, B, C]$ the parameters describing camera rotation¹ along the x, y and z axes respectively. Let f be the camera’s focal length and Z the relative scene depth at a pixel location (x, y) . In this setting, the motion vector \vec{v} due to camera motion is given by:

$$\begin{aligned} \vec{v} &= \vec{v}_r + \vec{v}_t = \begin{pmatrix} u_r \\ v_r \end{pmatrix} + \begin{pmatrix} u_t \\ v_t \end{pmatrix}, \\ &= \begin{pmatrix} \frac{A}{f}xy - Bf - \frac{B}{f}x^2 + Cy \\ Af + \frac{A}{f}y^2 - \frac{B}{f}xy - Cx \end{pmatrix} + \begin{pmatrix} -fU+xW \\ -fV+yW \\ Z \end{pmatrix}, \end{aligned} \quad (1)$$

where \vec{v}_r and \vec{v}_t represent motion field vectors corresponding to camera rotation and translation respectively. Equation (2)² highlights an important property, namely that the flow due to camera rotation is only determined by the camera rotation parameters and the camera’s focal length. *The flow due to camera rotation is independent of the scene depth.* One can subtract this rotational motion component at each pixel to obtain a rotation-compensated flow field.

3.1.2 Rotation-compensated motion field

As shown in the flow equation (2), the rotation-compensated flow field \vec{v}_t is determined by the translational camera motion $[U, V, W]$, and the scene depth

¹The rotation parameters are often referred to as pitch, yaw and roll.

²In fact, this equation is an approximation, and only holds if the rotation angles are small (Longuet-Higgins et al., 1980). To obtain the exact rotational flow field one has to transform the 2D image points to 3D using perspective projection equations, rotate the points according to the camera’s rotation in 3D, backproject them onto the 2D image plane, and then measure the displacement.

Z . It comprises all the relevant information about the scene geometry, unlike the rotational component \vec{v}_r , which is independent of the scene geometry. The magnitude of the rotation-compensated flow is inversely related to scene depth, i.e., regions further away from the camera have small translational flow magnitude, and those closer to the camera have larger magnitudes. *The direction of \vec{v}_t (flow angle) however does not depend upon the scene depth:*

$$\theta = \begin{cases} \arccos(xW - fU), & \text{if } (yW - fV) > 0, \\ 2\pi - \arccos(xW - fU), & \text{otherwise.} \end{cases} \quad (3)$$

Figure 4 pictures the computation of the flow angle θ at pixel locations (x, y) , leading to an angle field as shown on the right. Where as Figure 4 pictures the angle field of pure camera translation, Figure 2 shows an angle field of a scene with camera translation and object motion. Here, independently moving objects, can be observed as discontinuities in angle. The angle of the rotation-compensated flow alone is independent of the scene geometry, thus independently moving objects stand out due to their different direction.

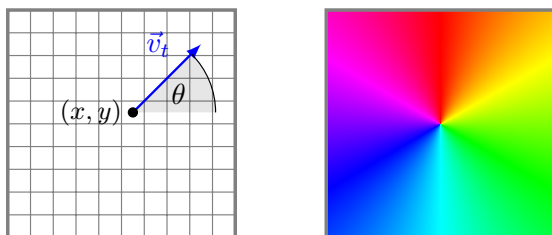


Fig. 4 Translational motion field vector. Left: motion field vector \vec{v}_t at pixel position (x, y) . Right: color coding of the angle field $\theta(x, y)$ at each pixel location for the case of camera translation along the optical axis, i.e. $[U, V, W] = [0, 0, 1]$.

3.2 The Right Spin: Camera Motion Estimation

To rectify the observed optical flow field for camera rotation, we require an accurate estimate for rotation. How can we obtain a good estimate of the camera rotation and the translational motion direction that together best explain the observed motion field? Towards finding an answer to this question, we derive a novel maximum likelihood approach that aims at finding the rotation $[A, B, C]$ such that the likelihood of the resulting translational flow field is maximized.

To this end, we derive a new flow likelihood function incorporating a model for the optical flow's noise as well as a prior distribution over the *inverse scene depth*.

In the following, we first introduce the new flow likelihood (Section 3.2.1). We then describe how camera motion parameters are estimated by maximizing this new likelihood function.

3.2.1 Likelihood of the translational motion field

Let \vec{o}_t be the observed translational flow vector, e.g., flow estimated with (Sun et al., 2018), at the pixel position (x, y) . Let the translational 3D motion direction of the camera $[U, V, W]$ be a unit vector. The three translational camera parameters $[U, V, W]$ and the pixel position (x, y) define the direction of a motion field vector on the image plane. As derived in (Bideau et al., 2018), the probability of observing \vec{o}_t at (x, y) given a motion direction $[U, V, W]$ is given by:

$$p(\vec{o}_t | U, V, W, x, y) = \int_0^\infty p(\vec{n}) p_r(r | U, V, W, x, y) dr, \quad (4)$$

where r denotes the magnitude of a motion field vector and \vec{n} is the optical flow's noise. This likelihood function depends on the distribution over the optical flow's noise $p(\vec{n})$ as well as the distribution over motion field magnitudes p_r . Figure 5 pictures the computation of $p(\vec{n})$. Modeling the probability distribution over flow magnitudes is challenging, since those depend on the camera's translational motion direction $[U, V, W]$, the pixel location as well as the scene depth at that location. (Bideau et al., 2018) model p_r by assuming that the motion field magnitude r is independent of the flow direction $[U, V, W]$. However this often does not lead to accurate estimations, especially in the case of strong z-motion (forward motion). Here, motion field magnitudes close to the focus of expansion are near zero and the motion vectors farther away from the focus of expansion show larger magnitudes, thus the motion field magnitude is clearly dependent upon the camera's motion direction.

Next, we present a new way of modeling the distribution over motion field magnitudes p_r that alleviates this problem.

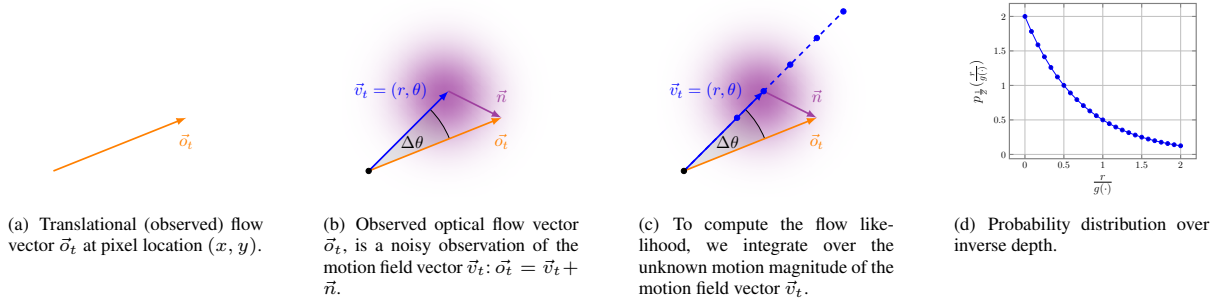


Fig. 5 Flow likelihood. (a)-(c): computation of the probability $p(\vec{n})$ at pixel location (x, y) . (d): probability distribution over inverse depth. The flow likelihood is maximal, when the observed flow vector \vec{o}_t and the motion field vector \vec{v}_t point into the same direction with similar magnitude.

Distributions over flow magnitudes

We express the motion field magnitudes in terms of inverse depth $\frac{1}{Z}$ and $g(\cdot)$, which is a function comprising all aspects of the magnitude that are *not* related to depth,

$$\begin{aligned} r &= \sqrt{u_t^2 + v_t^2}, \\ &= \frac{1}{Z} \cdot g(f, x, y, U, V, W). \end{aligned} \quad (5)$$

Given this reformulation of the magnitude r , we can determine the *induced distribution* over motion field magnitudes, given the distribution over inverse depths. We aim to compute $p_r(r | g(f, x, y, U, V, W))$ through $p_{\frac{1}{Z}}(\frac{1}{Z})$, which is the distribution over inverse depth. Using the relation between r and $g(\cdot)$ from Eq. 5, we can rewrite $p_r(r | g(\cdot))$ as follows

$$p_r(r | g(\cdot)) = \frac{p_{\frac{1}{Z}}(\frac{r}{g(\cdot)})}{g(\cdot)}. \quad (6)$$

This is effectively just a change of units. Expressing the distribution over flow magnitudes in terms of the distribution over inverse depth however brings a significant advantage. This formulation effectively factors motion direction (U, V, W) , focal length f and scene depth into the function $g(\cdot)$, and the distribution over depth can be modeled without relying on these dependencies that require making further approximations.

Flow likelihood

Following prior derivations, the flow likelihood function (Eq. 4) can be expressed by the distribution over

inverse depth, instead of flow magnitudes:

$$\begin{aligned} p(\vec{o}_t | U, V, W, x, y) &= \int_0^\infty p(\vec{n}) p_r(r | g(\cdot)) dr \\ &= \int_0^\infty p(\vec{n}) \frac{p_{\frac{1}{Z}}(\frac{r}{g(\cdot)})}{g(\cdot)} dr. \end{aligned} \quad (7)$$

The key advantage of this is that while flow magnitudes are not independent of the motion direction, the inverse depths are, and thus the model is more realistic.

3.2.2 Camera motion estimation via likelihood maximization

In Section 3.2.1, we have derived a new likelihood function of an observed optical flow vector \vec{o} . Our goal is now to find a camera rotation (A, B, C) and translational camera motion direction (U, V, W) , such that the flow likelihood is maximal or alternatively the negative log-likelihood is minimal. Recall \vec{o}_t is the observed translational flow vector after subtracting the flow \vec{v}_r due to camera rotation:

$$\vec{o}_t = \vec{o} - \vec{v}_r(A, B, C). \quad (8)$$

Given the rotation compensated flow, we minimize the negative log-likelihood as follows:

$$\begin{aligned} &A^*, B^*, C^*, U^*, V^*, W^* \\ &= \arg \min_{A, B, C, U, V, W} \sum -\log(p(\vec{o}_t | U, V, W, x, y)). \end{aligned} \quad (9)$$

Local minima are a concern when solving this optimization problem, especially in cases of noisy optical

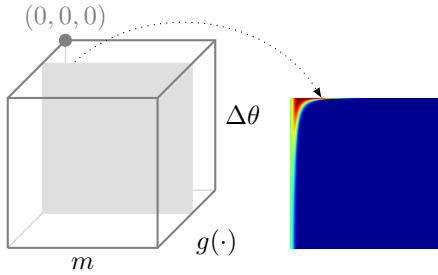


Fig. 6 Lookup table picturing flow likelihood values. Our new flow likelihood addresses the challenge of estimating the camera’s motion in the presence of noisy optical flow. The color *red* indicates high likelihood values, *dark blue* indicates low likelihood values. The lower the angle difference $\Delta\theta$ between the vectors \vec{o}_t and \vec{v}_t , the higher the likelihood. Note that for very small flow magnitudes m the flow likelihood is almost the same regardless $\Delta\theta$. This is an important consequence of our model, indicating the unreliability of the flow direction in case of near zero magnitudes.

flow, inaccurate estimates of independently moving objects or complex scene geometry. To reduce this risk, we initialize the optimization using three different starting points: (1) camera rotation and translation estimate of the previous frame, (2) camera rotation estimate weighted by depth estimate of the previous frame and the translation estimate of the previous frame, and (3) camera rotation estimate weighted by depth estimate of the previous frame and the translation estimate of the previous frame in the opposite direction. The first initialization is a good assumption if the camera motion is approximately constant. Initialization (2) and (3) incorporate depth information. The apparent motion of areas far away is mainly influenced by the camera’s rotation and not the camera’s translation (see Figure 7), thus knowing the depth helps to correctly disentangle flow due to camera rotation and flow due to translation.

During the optimization each pixel is weighted using learned, soft object motion masks of the previous frame, that evolve over time - thus the influence of moving objects is suppressed due to a low weight. The following Section describes how object motion masks are learned while pertaining important geometric information.

3.3 Object Motion Segmentation

We build our segmentation framework on an effective model for motion segmentation, that learns object motion patterns from optical flow and segments a flow field into static background and moving objects (Tokmakov et al., 2017a). Yet, this model does not incorporate any geometrical concepts. As discussed earlier

optical flow fields couple information about scene geometry as well as camera motion, making the judgment whether an object is moving challenging. By introducing a simple pre-processing step we show, that the complexity of optical flow patterns is dramatically reduced. Different from prior work, our network processes rotation compensated flow fields (angle + magnitude) to segment independently moving objects. Learning object motion based on pre-processed flow fields appears to be an easier task to learn. While our network architecture is similar to (Tokmakov et al., 2017a), we propose important modifications to the training procedure in the following.

3.3.1 Incorporating geometric information into training

The network follows the classical U-net architecture and is trained on estimated translational flow fields. During training, we first estimate optical flow using (Sun et al., 2018) on the FlyingThings3D data set (Mayer et al., 2016). The ground truth camera rotation is provided and subtracted from the estimated flow to obtain a rotation-compensated flow field. This flow field is input to our network. The input has a size of $h \times w \times 3$. The third dimension denotes the flow expressed in terms of angle (represented as a unit vector) and magnitude. A representation of the flow angle as unit vector instead of angles in degree avoids segmentation discontinuities at 0 degree (or 2π respectively). The normalized flow field and the flow’s magnitude are concatenated and form the input to our network. An interesting question for training a network with rotation-compensated optical flow is, whether it is worthwhile to incorporate the magnitude into the training procedure. On the one hand the flow magnitude can be a good indicator about the reliability of the flow angle (Bideau and Learned-Miller, 2016b), while on the other hand variation in larger magnitudes can be either due to variances in the scene depth or fast moving objects - thus including the magnitude might add rather misleading information. We take a closer look into this question as part of our ablation study in Section 4.2.

3.4 Implementation details

To find the camera rotation and translational motion direction that best explains the observed optical flow field, we derived a new flow likelihood function (Section 3.2). Details regarding parametrization are provided in the following.



Fig. 7 Flow, rotation compensated flow and the relative depth estimate. We show sample videos from the data set Complex Background (video sequences: traffic, forest) as well as two sample videos from the Davis data set (video sequence: parkour, goat). A comparison of (b) and (d) shows how motion at distant is dominated by camera rotation. After subtracting of the camera’s rotation the remaining flow magnitude in these areas is very small (light color). If the flow magnitude is small the motion direction is noisy. This can be seen in (e).

The probability of the flow noise $p(\vec{n})$ is modeled as a multivariate normal distribution $p(\vec{n}) \sim \mathcal{N}(\mu, \Sigma)$ and the inverse depth $p(\frac{1}{Z})$ as an exponential distribution $p(\frac{1}{Z}) \sim \text{Exp}(\lambda)$. The noise covariance Σ is assumed to be spherical and is measured using the ground truth flow of Sintel (Butler et al., 2012) and the corresponding noisy estimate Sun et al. (2018). We obtain $\Sigma = 16.5 \cdot 10^{-5}I$, where I is the identity matrix. λ is the rate parameter of the exponential distribution modeling the inverse depth, and is estimated using ground truth depths from Sintel. We measured $\lambda = 0.64$. The distribution over inverse depth can be seen in Figure 5(d).

For computational efficiency the integral in Eq. 7 is approximated using a discrete sum over motion

field magnitudes r . Flow likelihood values are pre-computed and stored in a lookup table for efficiency (see Figure 6).

4 Experiments

We begin with a brief description of data sets used for training and evaluation of our motion segmentation network. In Section 4.1, we evaluate our here presented motion segmentation approach on the widely used DAVIS (Perazzi et al., 2016) data set and MoCA (Lamdouar et al., 2020). Ground truth camera motion is not provided for these data sets, thus synthetic data - such as the FlyingThings3D data set (Mayer et al., 2016) and Sintel (Butler et al., 2012; Wulff et al., 2012) - are used for ablation studies.

These studies in particular focus on the analysis of different variants of our core network and the quality as well as the effect of rotation estimation via likelihood maximization.

DAVIS2016 (Densely Annotated Video Segmentation) (Perazzi et al., 2016) contains 50 video sequences in total with moving objects in various environments. A 30/20 training/validation split is provided. Our model is evaluated on the validation set. Ground truth segmentations of the most prominent moving object are provided for each frame. DAVIS has been widely used for general video segmentation as well as motion segmentation.

MoCA (Moving camouflaged animals) (Lamdouar et al., 2020) comprises a set of 141 videos depicting 67 different animals. The data set is split into three motion types describing the animals motion - *locomotion*, *static* and *deformation*. Following the procedure of (Lamdouar et al., 2021; Yang et al., 2019) we evaluate on the *locomotion* split, which forms the largest part of the dataset with 88 video sequences in total. Annotations are provided in form of bounding boxes. An evaluation script is provided by the authors of MoCA.

FT3D (FlyingThings3D) (Mayer et al., 2016) is a large optical flow data set, providing ground truth optical flow, RGB images, camera motion and depth. It is a synthetic data set showing random objects like chairs, tables, etc., flying in a 3D world along random trajectories. The data set is split into test and training sets.

Sintel (Butler et al., 2012; Wulff et al., 2012) is the de facto benchmark for optical flow algorithms, containing 23 video sequences with 20 to 50 frames each. These short video sequences are taken from an animated movie. The scenes are realistically simulated. Synthetic videos are available with ground truth optical flow, depth, camera motion and material segmentation.

4.1 Results

Our main framework consists of two steps (1) compensating the observed optical flow for camera rotation, and (2) segmenting the resulting optical flow in to static background and independently moving objects. Experiments presented here are based on the DAVIS (Perazzi et al., 2016) data set and the MoCA (Lamdouar et al., 2020) data set, that each raise a slightly different aspect onto the motion segmentation problem. Details are described in the following.

		LMP	TMM	Ours	Ours*
	Supervised	✓	✗	✓	✓
	RGB	✗	✗	✗	✗
	Flow	✓	✓	✓	✓
	Multi frame	✗	✗	✗	✗
\mathcal{J}	Mean \uparrow	58.4	40.1	59.7	62.5
	Recall \uparrow	67.3	34.3	69.6	73.8
	Decay \downarrow	5.6	15.2	4.3	3.8
\mathcal{F}	Mean \uparrow	58.4	39.6	59.5	61.1
	Recall \uparrow	66.0	15.4	66.4	69.9
	Decay \downarrow	7.9	12.7	5.4	5.6
\mathcal{T}	Mean \downarrow	87.8	51.3	74.5	83.4

Table 1 Motion segmentation: Comparison to other approaches using only motion cues on DAVIS (*train-val*), i.e., without any appearance. Ours refers to the variant of our model using only motion cues and no appearance terms and Ours* denotes a motion-only upper bound, which uses ground truth segmentation for camera motion estimation. Best viewed in color (**1st-best**, **2nd-best**).

DAVIS: Optical flow only

We compare our motion segmentation network with other methods that use optical flow as the only cue for segmentation. Table 1 shows these results on DAVIS. LMP (Tokmakov et al., 2017a) is a learning based approach trained on ground truth optical flow of FlyingThings3D. This approach relies on a similar network architecture, but does not incorporate an explicit model for modeling geometrical concepts, e.g. the scene geometry and camera motion. TMM (Bideau and Learned-Miller, 2016b), on the contrary, compensates flow for camera rotation and attempts to segment a video by assigning translational motion models to different image regions in a probabilistic fashion. The exclusive usage of translational motion models however quickly leads to oversegmentations and fails to capture more complex motion patterns. While combining geometrical concepts such as perspective projection together with learned motion patterns, our approach improves over both these motion segmentation methods. The segmentation performance is measured using the \mathcal{J} -Mean score. We achieve an \mathcal{J} -Mean score of 59.7. The next best performing method is LMP resulting in an \mathcal{J} -Mean score of 58.4. We compute an upper bound for our method (Ours* in Table 1) by masking out independently moving objects, with ground truth segments, for our camera motion estimation procedure. This masking procedure eliminates errors of our camera motion estimation due to ‘outliers’ in optical flow, such as moving objects.

		SegI	MG	CIS	COD	COSNet	MATNet	Ours	Ours+Temp
	Supervised	✗	✗	✗	✓	✓	✓	✓	✓
	RGB	✗	✗	✓	✓	✓	✓	✗	✗
	Flow	✓	✓	✓	✓	✗	✓	✓	✓
	Multi-frame	✓	✓	✓	✗	✗	✓	✗	✓
\mathcal{J}	Mean ↓	68.6	63.4	49.4	44.9	50.7	64.2	58.3	65.8
Success Rate	$\tau = 0.5$	77.2	74.2	55.6	41.4	58.8	71.2	64.5	72.7
	$\tau = 0.6$	71.7	65.4	33.0	53.4	67.0	46.3	58.0	65.2
	$\tau = 0.7$	62.3	52.4	32.9	23.5	45.7	59.9	49.8	53.8
	$\tau = 0.8$	46.4	35.1	17.6	14.0	33.7	49.2	36.2	38.9
	$\tau = 0.9$	25.5	14.7	3.0	5.9	16.7	24.6	16.6	17.3
	SR_{mean}	56.6	48.4	31.1	23.6	41.7	54.4	45.0	49.6

Table 2 Motion segmentation: Comparison to state-of-the-art motion segmentation methods on MoCA. Methods we compare against from left to right: (Lamdouar et al., 2021; Yang et al., 2021, 2019; Lamdouar et al., 2020; Lu et al., 2019b; Zhou et al., 2020). **Bold** indicates best among all methods, while **1st-best** and **2nd-best** represent the best and second best within the supervised methods. Best viewed in color.

MoCA: Optical flow only

Data sets like MoCA focus in particular on the segmentation of objects that can only be robustly recognized based on their unique motion. Where as most data sets for moving object segmentation combine several cues (motion and appearance) that are helpful for recognizing moving objects, this data set highlights the relevance of motion. Thus MoCA allows to evaluate the strengths of motion models in isolation. It is not surprising that appearance cues are rather weak in cases of camouflage, therefore methods based on RGB frames only (e.g. COSNet (Lu et al., 2019b)) show a weak performance in these settings (see Table 2). On more general data set like the DAVIS, those methods show a superior performance among all other methods and achieve a similar segmentation quality as MATNet (Zhou et al., 2020) (Table 3).

Our approach taking a single optical flow frame (compensated for camera rotation) as input, performs comparable to other supervised approaches. A simple post-processing step - convolution with a 3D Gaussian filter and frame-wise application of a dense CRF, eliminates temporal instabilities (Ours+Temp in Table 2). Among all methods SegI (Lamdouar et al., 2021) shows best results on MoCA, on DAVIS their performance falls rather short due to their lack of a strong appearance model. SegI combines multiple ConvNets where each of them encode a flow frame together with a transformer network without taking RGB frames into consideration. The model is trained on synthetically generated data, thus can be considered as unsupervised. In contrast, our approach was trained using

rotation compensated flow frames estimated from the synthetic dataset *FlyingThings3D*.

DAVIS: Optical flow + Appearance

Where our main contribution lays in a novel approach to learn to segment moving object based on optical flow only, we incorporate here appearance information similar to LVO (Tokmakov et al., 2017a) and compare to segmentation approaches that consider both - appearance as well as motion information (Table 3). Within the group of supervised approaches our approaches shows best performance in terms of mean/recall \mathcal{J} and \mathcal{F} . Where as ours and LVO integrate appearance cues in a similar manner, these approaches differ in the way how object motion cues are learned. LVO learns object motion patterns directly from optical flow, where as we first disentangle camera rotation and translation before segmenting independently objects. Ablation studies analyze the usefulness of this disentanglement in further detail. Within the unsupervised approaches ARP (Koh and Kim, 2017), which is a non learning based approach, reaches highest performance. Due to multiple iterations over the entire video this approach is computationally expensive as mentioned in (Yang et al., 2021). Among all methods MATNet reaches highest accuracy in terms of mean \mathcal{J} and \mathcal{F} . One reason might lay in their training strategy, which makes use of the DAVIS training set (indicated with ♣).

A qualitative comparison with the best performing methods is shown in Figure 9. Our results based on optical flow only and based on optical flow in combination with appearance are shown in the last two

		SFL	COSNet	MATNet	LMP+App	FSEG	LVO	Ours+App	CIS	ARP	SegI
	Supervised	♣	♣	♣	✓	✓	✓	✓	✗	✗	✗
	RGB	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗
	Flow	✓	✗	✓	✓	✓	✓	✓	✓	✓	✓
	Multi-frame	✗	✗	✓	✗	✗	✓	✓	✓	✓	✓
\mathcal{J}	Mean \uparrow	67.4	80.5	82.4	70.0	70.7	72.2	73.5	71.5	76.2	67.8
	Recall \uparrow	81.4	94.0	94.5	85.0	83.5	82.4	85.5	86.5	91.1	-
	Decay \downarrow	6.2	0.0	5.5	1.3	1.5	0.1	1.2	9.5	7.0	-
\mathcal{F}	Mean \uparrow	66.7	79.4	80.7	65.9	65.3	67.5	68.9	70.5	70.6	-
	Recall \uparrow	77.1	90.4	90.2	79.2	73.8	75.4	79.6	83.5	83.5	-
	Decay \downarrow	5.1	0.0	4.5	2.5	1.8	2.7	1.4	7.0	7.9	-

Table 3 Motion segmentation: Comparison to state-of-the-art motion segmentation methods on DAVIS2016. We group approaches according to their training strategy: supervised and trained on the DAVIS training split (♣), supervised and trained on other segmentation data sets (✓) and unsupervised methods (✗). Methods we compare against from left to right: (Cheng et al., 2017a; Lu et al., 2019b; Zhou et al., 2020; Tokmakov et al., 2017a; Jain et al., 2017; Tokmakov et al., 2017b; Yang et al., 2019; Koh and Kim, 2017; Lamdouar et al., 2021; Yang et al., 2021). **Bold** indicates best among all methods, while **1st-best** and **2nd-best** represent the best and second best within the supervised methods. Best viewed in color.

rows of this figure. These two rows in particular highlight the complementarity of motion and appearance cues. We miss the hiker’s foot when relying on motion alone (Ours), since it is not moving. However, while integrating motion with appearance, we segment the entire object accurately. ARP, the strongest method among unsupervised approaches, relies on segmenting the primary object(s) in a video and comes with a noticeable bias towards the object’s appearance. In many cases such a strong appearance model is advantageous, however can lead to erroneous segmentations in other cases. For example, it only segments a part of the car (Figure 9: 2nd column from the right), which moves from the darker (shadow) area to the brighter (sunny) region.), as it matches the primary object in appearance. Our method that extracts geometrical information from optical flow and integrates learned objectness cues is capable of overcoming these types of failure cases.

4.2 Ablation study

Network variants

We trained four variants of our motion segmentation network, with: (1) ground truth optical flow, (2) the ground truth flow after removing ground truth camera rotation, i.e., with rotation compensated-flow fields, (3) estimated optical flow field using PWC-Net (Sun et al., 2018), and (4) estimated ground truth flow compensated with ground truth camera rotation, i.e., estimated rotation compensated-flow field. Table 4 shows the analysis with these four variants. Training and testing with ground truth optical flow (original:

gt flow or compensated: gt t-flow) is significantly better than using estimated optical flow. Segmentation accuracy is about 20% higher on the FT3D test set for ground truth, compared to estimated optical flow. Training with rotation-compensated optical flow consistently leads to improved quality of the final segmentation, e.g., 90.68% vs. 93.23%, which supports the idea behind our method. Learning can be significantly simplified, if we are able to efficiently incorporate knowledge about physical concepts into the process of moving object segmentation. A direct comparison in terms of segmentation quality between using the original optical flow as input instead of the rotation-compensated optical flow is shown in Figure 8.

Training on flow angle only versus angle+magnitude

As discussed in Section 3.1, rotation-compensated flow comprises all the information about independent

trained with...	tested with...	angle+magnitude
gt flow	gt flow	90.68
gt t-flow	gt t-flow	93.23
PWC flow	PWC flow	77.18
PWC t-flow	PWC t-flow	78.69

Table 4 Ablation study: Network variants. We trained four networks using flow angle and magnitude with: the provided ground truth optical flow of FT3D (Mayer et al., 2016) (gt flow), ground truth optical flow after subtracting ground truth camera rotation (gt t-flow), estimated optical flow using (Sun et al., 2018) (PWC flow), and estimated optical flow after subtracting ground truth camera rotation (PWC t-flow). Segmentation accuracy is measured on the FT3D test set with intersection over union (IoU) scores.

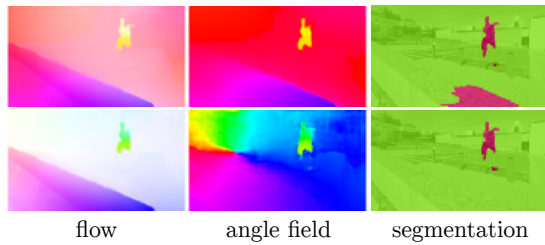


Fig. 8 Ablation study: Comparison of motion segmentation results based on the original and the rotation-compensated flow field. Top row: motion segmentation with the original flow field that includes camera rotation, translation and object motion. Bottom row: motion segmentation based on *rotation-compensated flow field*. Note that the angle field (middle) of the rotation-compensated flow is entirely depth independent. The angle field is fully determined by the translational camera motion and object motion. In this example one can observe a clear z-motion of the camera, which is shown by the rainbow pattern. The angle field of the original flow containing both camera rotation and translation is depth dependent (top row, middle image). This angle field clearly shows discontinuities in angle at the wall, which is due to significant changes in depth and not because of independent object motion.

object motion and the scene structure (depth). In this context, two interesting questions to tackle are: *how well can one extract information about independent object motion from the angle alone*, and *does including the flow magnitude (training the network on the full optical flow) improve motion segmentation?*. We show this analysis in Table 5, with further variants of our network. Using angle and magnitude together (angle+magn in the table) leads to the best performance. However, note that we achieve reasonable segmentation quality even when using the flow angle alone. The network trained on ground truth optical flow adapts very poorly to estimated optical flow, with the segmentation accuracy dropping from 93.23% to 24.44% for the angle+magn variant.

Rotation estimation via likelihood maximization

We show results on the Sintel data set (Table 6), and compare our new likelihood optimization procedure with (Bideau and Learned-Miller, 2016b). The ground truth optical flow and focal length is provided, so an accurate estimate of the camera’s rotation is possible. Our camera rotation estimation based on maximizing the flow likelihood shows consistently better results on the Sintel data set. More importantly, the performance gap gets significant when using estimated flow as input for camera motion estimation. Since our proposed optimization approach incorporates an explicit noise model, it is significantly more robust to noisy flow data.

trained with...	tested with...	angle	angle+magn
gt t-flow	gt t-flow	77.47	93.23
gt t-flow	PWC t-flow	24.06	24.44
PWC t-flow	PWC t-flow	77.79	78.69

Table 5 Ablation study: Training with angle vs angle and magnitude. We trained four variants of our segmentation network with: (1) angle of the rotation-compensated flow of FT3D, (2) angle and magnitude of the rotation-compensated flow (angle+magn), (3) angle of the estimated rotation-compensated flow, and (4) angle and magnitude of the estimated rotation-compensated flow. We show consistently better performance by including magnitude. The performance is the worst when the network is trained on the angle of the rotation-compensated ground truth flow. Here, the noise in angle leads to a very significant drop on estimated optical flow data. Segmentation accuracy is measured on the FT3D test set with intersection over union (IoU).

	TMM	Ours
gt-flow	0.08 / 0.22 / 0.02	0.03 / 0.06 / 0.01
PWC-flow	0.13 / 0.34 / 0.04	0.05 / 0.11 / 0.03

Table 6 Ablation study: Camera rotation estimation. Avg. yaw/pitch/roll error in degrees between two consecutive frames. *gt-flow*, *PWC-flow*: To evaluate rotation estimation we used ground-truth segmentation masks to weight the optim. loss. Thus, errors in the segmentation procedure are not propagated throughout the video.

Declarations

Funding

This work was supported in part by the ANR grant AVENUE (ANR-18-CE23-0011) and Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany’s Excellence Strategy – EXC 2002/1 “Science of Intelligence” – project number 390523135.

Conflicts of interest/competing interests

The authors have no conflicts of interest to declare that are relevant to the content of this article.

References

- Bideau, P. and E. Learned-Miller. 2016a. A detailed rubric for motion segmentation. *arXiv preprint arXiv:1610.10033*.
- Bideau, P. and E. Learned-Miller 2016b. It’s moving! A probabilistic model for causal motion segmentation in moving camera videos. In *Proceedings of*



Fig. 9 Qualitative segmentation results. Qualitative segmentation results on the DAVIS data set, showing a comparison with three other best performing methods. Ours-final denotes our complete method and Ours the variant based on motion cues alone.

the European Conference on Computer Vision.

Bideau, P., R.R. Menon, and E. Learned-Miller 2018, September. Moa-net: Self-supervised motion segmentation. In *Proceedings of the European Conference on Computer Vision Workshops.*

Bideau, P., A. RoyChowdhury, R.R. Menon, and E. Learned-Miller 2018. The best of both worlds: Combining CNNs and geometric constraints for hierarchical motion segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition.*

Brox, T. and J. Malik 2010. Object segmentation by long term analysis of point trajectories. In *Proceedings of the European Conference on Computer Vision.*

Butler, D.J., J. Wulff, G.B. Stanley, and M.J. Black 2012. A naturalistic open source movie for optical flow evaluation. In *Proceedings of the European Conference on Computer Vision.*

Cheng, J., Y.H. Tsai, S. Wang, and M.H. Yang 2017a. Segflow: Joint learning for video object segmentation and optical flow. In *Proceedings of the IEEE International Conference on Computer Vision.*

Cheng, J., Y.H. Tsai, S. Wang, and M.H. Yang 2017b. SegFlow: Joint learning for video object segmentation and optical flow. In *Proceedings of the IEEE International Conference on Computer Vision.*

Dave, A., P. Tokmakov, and D. Ramanan 2019. Towards segmenting anything that moves. In *Proceedings of the IEEE International Conference on Computer Vision Workshops.*

Faktor, A. and M. Irani 2014. Video segmentation by non-local consensus voting. In *Proceedings of the British Machine Vision Conference.*

Fragkiadaki, K., G. Zhang, and J. Shi 2012. Video segmentation by tracing discontinuities in a trajectory embedding. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition.*

Gordon, A., H. Li, R. Jonschkowski, and A. Angelova 2019. Depth from videos in the wild: Unsupervised monocular depth learning from unknown cameras. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 8977–8986.

Horn, B.K. 1999. Projective geometry considered harmful. *Unpublished Memo*.

- Ilg, E., N. Mayer, T. Saikia, M. Keuper, A. Dosovitskiy, and T. Brox 2017. Flownet 2.0: Evolution of optical flow estimation with deep networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*.
- Irani, M. and P. Anandan. 1998. A unified approach to moving object detection in 2D and 3D scenes. *IEEE Transactions Pattern Analysis and Machine Intelligence* 20(6): 577–589 .
- Jain, S., B. Xiong, and K. Grauman 2017. Fusion-Seg: Learning to combine motion and appearance for fully automatic segmentation of generic objects in videos. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*.
- Jin, Y., L. Tao, H. Di, N.I. Rao, and G. Xu 2008. Background modeling from a free-moving camera by multi-layer homography algorithm. In *Proceedings of the IEEE International Conference on Image Processing*.
- Ke, Q. and T. Kanade 2002. A robust subspace approach to layer extraction. In *Workshop on Motion and Video Computing*. IEEE.
- Keuper, M. 2017. Higher-order minimum cost lifted multicuts for motion segmentation. In *Proceedings of the IEEE International Conference on Computer Vision*.
- Keuper, M., B. Andres, and T. Brox 2015. Motion trajectory segmentation via minimum cost multicuts. In *Proceedings of the IEEE International Conference on Computer Vision*.
- Koh, Y.J. and C.S. Kim 2017. Primary object segmentation in videos based on region augmentation and reduction. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*.
- Lai, Z., E. Lu, and W. Xie 2020. Mast: A memory-augmented self-supervised tracker. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*.
- Lamdouar, H., W. Xie, and A. Zisserman 2021. Segmenting invisible moving objects. In *British Machine Vision Conference*.
- Lamdouar, H., C. Yang, W. Xie, and A. Zisserman 2020. Betrayed by motion: Camouflaged object discovery via motion segmentation. In *Proceedings of the Asian Conference on Computer Vision*.
- Lezama, J., K. Alahari, J. Sivic, and I. Laptev 2011. Track to the future: Spatio-temporal video segmentation with long-range motion cues. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*.
- Longuet-Higgins, H.C., K. Prazdny, et al. 1980. The interpretation of a moving retinal image. *Proceedings of the Royal Society of London. Series B, Biological sciences* 208(1173): 385–397 .
- Lu, X., W. Wang, C. Ma, J. Shen, L. Shao, and F. Porikli 2019a. See more, know more: Unsupervised video object segmentation with co-attention siamese networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*.
- Lu, X., W. Wang, C. Ma, J. Shen, L. Shao, and F. Porikli 2019b. See more, know more: Unsupervised video object segmentation with co-attention siamese networks. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Mahadevan, S., A. Athar, A. Ošep, S. Hennen, L. Leal-Taixé, and B. Leibe 2020. Making a case for 3d convolutions for object segmentation in videos. In *Proceedings of the British Machine Vision Conference*.
- Mayer, N., E. Ilg, P. Hausser, P. Fischer, D. Cremers, A. Dosovitskiy, and T. Brox 2016. A large dataset to train convolutional networks for disparity, optical flow, and scene flow estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*.
- Narayana, M., A. Hanson, and E. Learned-Miller 2013. Coherent motion segmentation in moving camera videos using optical flow orientations. In *Proceedings of the IEEE International Conference on Computer Vision*.
- Ochs, P. and T. Brox 2011. Object segmentation in video: A hierarchical variational approach for turning point trajectories into dense regions. In *Proceedings of the IEEE International Conference on*

- Computer Vision*. IEEE.
- Ogale, A.S., C. Fermüller, and Y. Aloimonos. 2005. Motion segmentation using occlusions. *IEEE Transactions Pattern Analysis and Machine Intelligence* 27(6): 988–992 .
- Papazoglou, A. and V. Ferrari 2013. Fast object segmentation in unconstrained video. In *Proceedings of the IEEE International Conference on Computer Vision*.
- Perazzi, F., J. Pont-Tuset, B. McWilliams, L. Van Gool, M. Gross, and A. Sorkine-Hornung 2016. A benchmark dataset and evaluation methodology for video object segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*.
- Pirenne, M.H. 1952. The scientific basis of Leonardo da Vinci’s theory of perspective. *The British Journal for the Philosophy of Science* 3(10): 169–185 .
- Ranjan, A., V. Jampani, L. Balles, K. Kim, D. Sun, J. Wulff, and M.J. Black 2019. Competitive collaboration: Joint unsupervised learning of depth, camera motion, optical flow and motion segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*.
- Shen, J., J. Peng, and L. Shao. 2018. Submodular trajectories for better motion segmentation in videos. *IEEE Transactions on Image Processing* 27(6): 2688–2700 .
- Sun, D., S. Roth, and M.J. Black 2010. Secrets of optical flow estimation and their principles. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*.
- Sun, D., X. Yang, M.Y. Liu, and J. Kautz 2018. PWC-Net: CNNs for optical flow using pyramid, warping, and cost volume. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*.
- Taylor, B., V. Karasev, and S. Soatto 2015. Causal video object segmentation from persistence of occlusions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*.
- Tokmakov, P., K. Alahari, and C. Schmid 2017a. Learning motion patterns in videos. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*.
- Tokmakov, P., K. Alahari, and C. Schmid 2017b. Learning video object segmentation with visual memory. In *Proceedings of the IEEE International Conference on Computer Vision*.
- Torr, P.H. 1998. Geometric motion segmentation and model selection. *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences* 356(1740): 1321–1340 .
- Tung, H.Y.F., R. Cheng, and K. Fragkiadaki 2019. Learning spatial common sense with geometry-aware recurrent networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*.
- Vertens, J., A. Valada, and W. Burgard 2017. SMSnet: Semantic motion segmentation using deep convolutional neural networks. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*.
- Vidal, R. and Y. Ma 2004. A unified algebraic approach to 2-D and 3-D motion segmentation. In *Proceedings of the European Conference on Computer Vision*.
- Vidal, R., S. Soatto, Y. Ma, and S. Sastry 2002. Segmentation of dynamic scenes from the multibody fundamental matrix. In *Proceedings of the European Conference on Computer Vision Workshops*.
- Walls, G. 1962. The evolutionary history of eye movements. *Vision Research* 2(1-4): 69–80 .
- Wang, J.Y. and E.H. Adelson. 1994. Representing moving images with layers. *IEEE Trans. Image Processing* 3(5): 625–638 .
- Wulff, J., D.J. Butler, G.B. Stanley, and M.J. Black 2012. Lessons and insights from creating a synthetic optical flow benchmark. In *Proceedings of the European Conference on Computer Vision Workshops*.
- Xiao, J. and M. Shah. 2005. Motion layer extraction in the presence of occlusion using graph cuts.

IEEE Transactions Pattern Analysis and Machine Intelligence 27(10): 1644–1659 .

Xu, X., L. Fah Cheong, and Z. Li 2018. Motion segmentation by exploiting complementary geometric models. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*.

Yan, J. and M. Pollefeys 2006. A general framework for motion segmentation: Independent, articulated, rigid, non-rigid, degenerate and non-degenerate. In *Proceedings of the European Conference on Computer Vision*.

Yang, C., H. Lamdouar, E. Lu, A. Zisserman, and W. Xie 2021. Self-supervised video object segmentation by motion grouping. In *Proceedings of the IEEE International Conference on Computer Vision*.

Yang, G. and D. Ramanan 2021. Learning to segment rigid motions from two frames. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*.

Yang, Y., A. Loquercio, D. Scaramuzza, and S. Soatto 2019. Unsupervised moving object detection via contextual information separation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*.

Zamalieva, D. and A. Yilmaz. 2014. Background subtraction for the moving camera: A geometric approach. *Computer Vision and Image Understanding* 127: 73–85 .

Zhang, G., J. Jia, W. Xiong, T.T. Wong, P.A. Heng, and H. Bao 2007. Moving object extraction with a hand-held camera. In *Proceedings of the IEEE International Conference on Computer Vision*.

Zhou, T., J. Li, S. Wang, R. Tao, and J. Shen. 2020. Matnet: Motion-attentive transition network for zero-shot video object segmentation. *IEEE Transactions on Image Processing* 29: 8326–8338 .