

Dense Monocular Motion Segmentation Using Optical Flow and Pseudo Depth Map: A Zero-Shot Approach

Yuxiang Huang
Systems Design Engineering
University of Waterloo
Waterloo, Canada
yuxiang.huang@uwaterloo.ca

Yuhao Chen
Systems Design Engineering
University of Waterloo
Waterloo, Canada
yuhao.chen1@uwaterloo.ca

John Zelek
Systems Design Engineering
University of Waterloo
Waterloo, Canada
jzelek@uwaterloo.ca

Abstract—Motion segmentation from a single moving camera presents a significant challenge in the field of computer vision. This challenge is compounded by the unknown camera movements and the lack of depth information of the scene. While deep learning has shown impressive capabilities in addressing these issues, supervised models require extensive training on massive annotated datasets, and unsupervised models also require training on large volumes of unannotated data, presenting significant barriers for both. In contrast, traditional methods based on optical flow do not require training data, however, they often fail to capture object-level information, leading to over-segmentation or under-segmentation. In addition, they also struggle in complex scenes with substantial depth variations and non-rigid motion, due to the overreliance of optical flow. To overcome these challenges, we propose an innovative hybrid approach that leverages the advantages of both deep learning methods and traditional optical flow based methods to perform dense motion segmentation without requiring any training. Our method initiates by automatically generating object proposals for each frame using foundation models. These proposals are then clustered into distinct motion groups using both optical flow and relative depth maps as motion cues. The integration of depth maps derived from state-of-the-art monocular depth estimation models significantly enhances the motion cues provided by optical flow, particularly in handling motion parallax issues. Our method is evaluated on the DAVIS-Moving and YTVOS-Moving datasets, and the results demonstrate that our method outperforms the best unsupervised method and closely matches with the state-of-the-art supervised methods.

Index Terms—Monocular Motion Segmentation; Zero-Shot Learning

I. INTRODUCTION

Dense motion segmentation aims to partition a video frame into separate areas characterized by consistent motion. The ability to identify and segment moving objects with a moving camera is crucial for various downstream applications, including autonomous navigation, robotics, simultaneous localization and mapping (SLAM), and holistic scene understanding. In dynamic environments, the video camera moves at an unknown speed relative to its environment, posing significant challenges for motion segmentation methods, such as motion parallax and motion degeneracy [1].

Existing techniques for dense monocular motion segmentation can be divided into two main categories: traditional methods which heavily rely on the brightness constancy constraint (usually in the form of optical flow) for the motion cue, and deep learning methods that can learn multiple motion and appearance cues from the video sequence through training. However, both types of methods have inherent limitations. Deep learning based approaches can produce impressive results on complex scenes where there are multiple moving objects, different types of motions and significant depth variations, but the state-of-the-art deep learning methods require end-to-end training with a significant amount of supervision, which is computation intensive and does not generalize well to different scenes [2]–[7]. Conversely, most traditional methods rely heavily on the optical flow or the brightness constancy constraint, which significantly limits their performance in complex scenes with significant depth variations or non-rigid motions [8]–[12]. Moreover, traditional methods are not able to accurately delineate moving objects as a whole [12], [13], especially when there are multiple moving components of the same object, due to their inability to learn the high-level appearance information of the objects.

To illustrate why the overreliance on optical flow limits the performance of motion segmentation methods, figure 1 shows an example of why optical flow by itself is insufficient if used as the only motion cue in motion segmentation. In this image sequence, the camera is moving forward and the cyclist is moving towards the camera. This particular image sequence has significant motion parallax since most objects in the scene are at different depths, which makes it almost impossible to tell which part of the scene is moving simply by looking at the optical flow field. This is because optical flow vectors are 2D projections of 3D velocity vectors and such projection is determined by both the depth and the screw motion of the object [14]. Therefore, when the optical flow is used as the only motion cue, the proposed method is not able to correctly segment the cyclist who is the only moving part in the scene. However, when the monocular depth map is jointly used with the optical field as motion cues, the method is able to produce




Fig. 1: Motion segmentation results of the proposed method by using only optical flow vs. using both optical flow and relative depth as the motion cue. (a) is a frame from a video sequence. (b) is the object proposal generated for this input frame. (c) and (d) are the optical flow mask and the relative monocular depth map generated by the off-the-shelf deep learning models. (e) and (f) show the motion segmentation results of our method by using only optical flow and optical flow + relative depth. In this case, optical flow alone is insufficient to segment the moving object due to motion parallax as well as forward motion.

the correct motion segmentation results.

To overcome these limitations of the existing methods, we propose a zero-shot monocular motion segmentation pipeline that is able to generate high-quality motion segmentation results without any training. We leverage the strong zero-shot capability of the computer vision foundation models to automatically recognize, detect, segment and track potential objects in the scene. Using this method, we can generate a high-quality object proposal for every frame in the image sequence. We then compute object-specific optical flow and depth map for each object on a per-frame basis. The depth map is generated by an off-the-shelf monocular estimation network, so it is only a pseudo depth map indicating relative depth. By analyzing how well each object’s optical flow mask fits on every other objects given their relative depth, we are able to derive a motion affinity matrix representing the pairwise motion similarities between all pairs of objects in the object proposal. We then apply spectral clustering on the motion affinity matrix to cluster different objects into different motion groups and obtain the final motion segmentation.

Our method was evaluated on two widely recognized motion segmentation benchmark datasets: DAVIS-Moving and YTVOS-Moving. Although our method does not require any training, experiments show that is closely matched with the state-of-the-art supervised methods and is superior to the best unsupervised method. In summary, the key contributions are

as follows:

- 1) We build the first zero-shot motion segmentation pipeline to achieve high-quality dense monocular motion segmentation without requiring any training.
- 2) We show the effectiveness of using relative depth cue as a complement motion cue in improving optical flow based motion segmentation algorithms.

II. RELATED WORK

Dense monocular motion segmentation methods can generally be divided into two groups: (1) Intensity based methods [8]–[10], [13], [15], [16] and (2) deep learning based methods [2]–[5], [7], [17]–[21]. Besides these two groups of methods which aim to generate a dense segmentation mask, there are also other motion segmentation methods that perform motion segmentation by clustering pre-computed point trajectories into different motion groups [22]–[30]. These methods do not belong to dense motion segmentation, but since they are relevant to our proposed method, we will briefly introduce them too. Additionally, it is important to differentiate between motion segmentation and video object segmentation (VOS) [31]. The goal of VOS is to segment only the moving objects in the foreground, while motion segmentation focuses on segmenting any object that is moving independently, regardless of whether it is in the foreground or not.

A. Intensity Based Methods

Intensity based methods are traditional methods based on the image brightness constancy constraint, which presumes the pixel intensity of an image stays constant over a short period of time. Intensity based methods can be further categorized into indirect and direct methods. Indirect methods [8]–[10], [13] use pixel-wise correspondences as input and generate a pixel-wise segmentation mask that delineates different motion groups. Such pixel correspondences are usually obtained from optical flow, which is based on the brightness constancy constraint. In contrast, direct methods [15], [32]–[34] directly take a pair of images as input and combine the two processes of optimizing for the brightness constancy constraint and estimating the motion models together. Most recent works on intensity based methods use optical flow based indirect methods, possibly due to the fast advance in optical flow estimation [35], [36]. Intensity based methods typically use causal inference or iterative optimization techniques to estimate the motion regions and motion models simultaneously [8]–[10], [13].

Intensity-based methods work well on simple scenes where the object motion and scene structure are simple, but will fail on more complex objects or scenes. For example, the motion of a walking human may not be modeled by a single optical flow mask since it may contain multiple moving parts (legs, arms and torso). In this case, these methods can have problems like over-segmentation or under-segmentation, where they segment different human parts as different motions, or only segment part of the human as the moving object. In addition, if the scene exhibits significant depth variation, these methods often struggle to determine whether a part of the image is moving independently or is simply located at a different depth compared to its surroundings.

B. Deep Learning Based Methods

Most deep learning methods take a sequence of image frames and sometimes also the precomputed motion cues such as the optical flow field or the monocular depth maps of these images, and produce a dense segmentation map in an end-to-end manner. Earlier methods are often only able to perform binary motion segmentation [3], [4], [17], [18], [20], but more recent methods have made significant progress and are able to achieve promising results in multi-label motion segmentation [2], [5]–[7], [21], [37].

Many deep learning based methods adopt a fully supervised approach [2], [3], [5], [17]. These methods typically train a CNN-based encoder-decoder network to perform end-to-end learning, which is computation-intensive. Their network architecture usually have the following components: (1) a module to extract the motion information from consecutive image frames, (2) a module to extract appearance information from the same sequence of frames, (3) a module to fuse the appearance and motion information together, and (4) a decoder to generate the final segmentation. These methods perform very well on scenes similar to the datasets they are trained on, but cannot scale well to unseen environment where there are

different motion patterns or object classes. Moreover, the data collection and training process are very time-consuming and computation intensive, which make them not an ideal method.

Aside from supervised methods, some methods also use semi-supervised, self-supervised and unsupervised approaches. In [19], the authors extended their previous work [13] by proposing a self-supervised approach to train a neural network to perform motion segmentation on synthetic angle fields, given that most optical flows can be reduced to rotation-compensated angle fields. In [38], the authors proposed an unsupervised learning method to solve multi-label motion segmentation problem by training a neural network to mimic the motion segmentation results from the Expectation-Maximization (EM) algorithm. However, these two methods purely rely on optical flow for motion information and thereby inheriting its limitations. To alleviate this problem, [37] proposes to train image segmentation and motion segmentation models together using both optical flow and raw video frames as inputs due to the fact that motion and appearance cues are usually highly related in practice. The unsupervised training is done in a very similar way as [38] using the EM-algorithm.

C. Sparse Correspondence Based Methods

Unlike Intensity based methods or deep learning methods, sparse correspondence based methods cannot produce dense segmentation masks of different moving objects. Instead, they output clusters of predefined keypoints corresponding to different motion groups instead of dense segmentation masks. These methods can be further categorized into two-frame based methods and multi-frame based methods. Two-frame based methods [22], [23], [39], [40] typically determine motion parameters using iterative optimization methods. This involves identifying a specific number of geometric models, such as homography, based on a set of corresponding keypoints, aiming to minimize an energy function representing the overall quality of the corresponding keypoints clustering. Unlike two-frame based methods, multi-frame based methods [25]–[30], [41]–[45] usually establish point correspondences over multiple frames using an optical flow based point tracker. Noisy, occluded and unwanted points are often manually removed to produce a sparse set of completely noise-free point trajectories. Multi-frame based methods have proven to be superior than two-frame based methods due to their ability to analyze the motion data in a longer time window using various geometric models and spatio-temporal similarities. Moreover, unlike two-frame based methods which only rely on epipolar geometry to detect different motions, multi-frame based methods are able to combine different geometric models together and achieve impressive results in challenging scenes with complex motions.

III. MOTION SEGMENTATION PIPELINE

We introduce a novel monocular dense motion segmentation pipeline that performs . The method first automatically extracts high-quality object proposals using the computer vision foundation models, and then clusters these proposed objects into




Fig. 2: Diagram of our proposed motion segmentation method. Given an image sequence, the proposal extraction module automatically identifies, segments and tracks all common objects through the whole sequence to generate an object proposal for each frame. Meanwhile, the motion cue generation module generates optical flow masks and monocular depth maps using PWC-Net and DINOv2. object-specific optical flow and depth maps are obtained by combining object proposals with optical flow and monocular depth maps. Given these object-specific motion cues, pairwise object similarity scores are computed to construct the motion similarity matrix. Finally, spectral clustering is used to cluster each object into its motion group.

distinct motion groups according to the motion cues provided by the object-specific optical flow and relative depth maps. Figure 2 shows a diagram of our motion segmentation pipeline.

A. Automatic Object Proposal Extraction

To automatically identify all moving objects within a video, the first task is to recognize and detect each common object in the video and track their trajectories throughout the video. We leverage the recently proposed computer vision foundational models in object recognition (RAM) [46], detection (Grounding DINO) [47], and segmentation (SAM) [48], alongside a state-of-the-art object tracking model (DeAOT) [49]. This preprocessing pipeline is based on the Segment and Track Anything (SAMTrack) [50] framework, which integrates the aforementioned models into a unified object segmentation and tracking system. SAMTrack generates dense object tracking masks based on a user-defined textual prompt specifying the desired objects to be tracked. To automate our system and eliminate the need for manual text prompts, we incorporate RAM at the beginning of our pipeline to automatically identify common objects in the initial video frame.

In summary, our complete preprocessing pipeline involves using RAM to identify common objects in the first frame, using the output from RAM as a textual prompt for the Grounding DINO model to obtain object bounding boxes, and then using these bounding boxes with SAM to generate an instance segmentation mask of the first frame. Non-max

suppression is applied to remove overlapping objects with an IoU score greater than 0.5 or with a mask area exceeding half the image size. Finally, The DeAOT tracker is used to follow each object's mask throughout the video.

To account for potential new objects entering the scene midway through the video, we divide the video into multiple parts consisting of equal numbers of frames and apply the full preprocessing pipeline to each part individually. The specific length of the each individual part can vary, but it is typically more beneficial to set a smaller number in more dynamic videos where more objects enter the scene midway.

B. Optical Flow and Relative Depth as Object-Specific Motion Cues

In order to determine if a set of objects have the same motion, we need to analyze object-specific motion cues for each object. More specifically, we calculate a dense optical flow mask and a monocular relative depth map for every video frame as the motion cues. The extraction of optical flow masks and relative depth maps are conducted using PWC-Net [51] and DINOv2 [52] respectively. Both models are the state-of-the-art models in their respective domains.

Relying solely on optical flow for motion segmentation is inadequate, as it cannot effectively distinguish between different motions and different depths. This limitation becomes evident when a camera moves, causing two stationary objects at different depths to appear as if they are moving differently

due to motion parallax. To address this limitation, it is essential to integrate depth information with optical flow, enhancing its ability to accurately analyze motion. In the following section, we present a parametric model that combines optical flow and depth data. We will demonstrate how this model can be used to compute the three-dimensional screw motions of objects, thus enabling the differentiation of objects based on their motions. This parametric model offers a robust theoretical framework for understanding complex motions in dynamic scenes.

C. Parametric Motion Model

We propose a motion model fitting algorithm that uses a parametric model derived from optical flow and depth data to represent the motions of individual objects throughout the video. This parametric motion model incorporates a revised version of the model equations introduced by Longuet-Higgins and Pruzdny [53], which can be used to compute the instantaneous screw motion of rigid objects at arbitrary depths. The original Longuet-Higgins and Pruzdny model equations establish a relationship between the optical flow, the instantaneous screw motion of rigid objects and the depths of individual pixels as follows:

$$\begin{aligned} u &= -\frac{xy}{f}\omega_1 + \frac{f^2 + x^2}{f}\omega_2 - y\omega_3 + \frac{f\tau_1 - x\tau_3}{z} \\ v &= -\frac{f^2 + y^2}{f}\omega_1 + \frac{xy}{f}\omega_2 + x\omega_3 + \frac{f\tau_2 - y\tau_3}{z} \end{aligned} \quad (1)$$

Here, u and v denote the optical flow vectors along the x and y axes respectively, z denotes pixel depth, f stands for the camera's focal length, and $\tau_1, \tau_2, \tau_3, \omega_1, \omega_2, \omega_3$ symbolize the object's translational and rotational movements. Nonetheless, the absolute depth of each pixel is often unknown in practice, making the complete utilization of this model for computing object motion unfeasible. To overcome this limitation, existing methods often use a parametric equation to infer object motion directly from the optical flow without knowing depth. For instance, [21] applies a segmented parametric formula with 12 parameters to precisely align with the optical flow field:

$$\begin{aligned} u &= a + bx + cy + dx^2 + exy + fy^2 \\ v &= g + hx + iy + jx^2 + kxy + ly^2 \end{aligned} \quad (2)$$

However, such a motion model lacks theoretical accuracy and fails to accommodate scenes with significant depth variations. Other research [13], [16] employs a simpler parametric motion equation to estimate the rotation-compensated optical flow angle field, albeit requiring known camera intrinsic parameters, which is impractical. To establish a motion model that is both theoretically robust and independent of camera intrinsics, we propose to linearize the Longuet-Higgins and Pruzdny equations using the monocular depth map produced from DINOv2. With the relative depth of each pixel known, we can reformulate the original equations into the following linear parametric form:

$$\begin{aligned} u &= a + b\frac{1}{z} - c\frac{x}{z} - dy + ex^2 - fxy \\ v &= g + h\frac{1}{z} - c\frac{y}{z} - dx + exy + fy^2 \end{aligned} \quad (3)$$

This set of linearized equations is more robust to motion parallax than (2) where the depth value is encoded in multiple unknown parameters and need to be approximated together with the screw motions. Using relative depth is sufficient for our goal, which is to distinguish different motions instead of computing the absolute values of the screw motion parameters.

D. Pairwise Motion Similarity Matrix

Once all optical flow motion models are obtained, each object will have a parametric motion model for every pair of frames. By fitting every object's optical flow vectors and depth map on its parametric motion model for the same frame pair, we can compute the residuals of each object to the motion models of all other objects. The residuals are computed with the mean squared error. Given N objects proposed for the scene in total, the motion residual vector for the i -th object at frame pair m are derived as follows:

$$\mathbf{e}_i^m = [e_{i,1}^m, e_{i,2}^m, e_{i,3}^m, \dots, e_{i,N}^m],$$

where $e_{i,n}^m$ represents the mean residual calculated by fitting the motion model of object i to the optical flow and depth of object n for frames m and $m+1$. A motion similarity matrix is then constructed to encode the motion similarity scores across all pairs of objects. This is achieved using the ordered residual kernel (ORK) [54]. To do so, the residual values in each residual vector are first sorted in ascending order. A threshold t is set to select the lowest t -th residuals as inliers. An binary inlier vector $\mathbf{v}_i = \{0, 1\}^N$ is then obtained for each object, where N is the total number of objects. The pairwise motion similarity between objects i and j can be calculated as $\mathbf{d}_{ij} = \mathbf{v}_i^T \mathbf{v}_j$, representing the frequency at which these two objects occur as each other's motion inliers. The ORK selects a certain number of inliers from each object's motion vectors instead selecting inliers below a threshold, making it more robust to different scenes and motions. After the motion similarity matrix is constructed, each motion similarity score \mathbf{d}_{ij} is normalized by dividing it with the number of frames that objects i and j have in common. This normalization step helps eliminate the weighting bias caused by incomplete trajectories.

E. Clustering Objects into Distinct Motion Groups

We apply row normalization [55] to the constructed motion similarity matrix and use spectral clustering to cluster objects into different motion groups. Given a predefined number of motion groups, spectral clustering cluster objects with high motion similarity values into the same motion group. Spectral clustering is a widely adopted technique in sparse correspondence based motion segmentation methods. It has shown effectiveness in clustering different point trajectories



Fig. 3: Qualitative comparison with the state-of-the-art unsupervised dense motion segmentation method [21] on DAVIS-Moving (column 1 - 3) and YTVOS-Moving (column 4 - 6) datasets. First row: Original video frame. Second row: Motion segmentation results produced by [21]. Third row: Motion segmentation results of our method. Last row: Ground truth

into different motion groups [29], [41], [56]. Therefore, we use it to cluster different objects in the object proposal into different motion groups given the motion similarity matrix.

IV. EXPERIMENTS

We conduct experiments using two widely recognized datasets: DAVIS-Moving and YTVOS-Moving. This section provides an overview of these datasets and the evaluation metrics, and compares our method with state-of-the-art approaches. Additionally, we present an ablation study to analyze the contributions of optical flow and depth information in enhancing motion segmentation results compared to the baseline.

A. Datasets and Evaluation Metrics

DAVIS-Moving and YTVOS-Moving [5] are the most recent benchmarks for generic instance motion segmentation. They are subsets of the DAVIS-17 dataset [57] and the YTVOS dataset [58] respectively. On the contrary to the original DAVIS and YTVOS datasets which focus on video object segmentation and only label moving objects in the foreground, the DAVIS-Moving and YTVOS-Moving datasets contain only videos where all moving objects are labeled. These datasets are particularly challenging due to occlusions, non-rigid motions and the diversity of object classes. We evaluate our method using the precision (Pu), recall (Ru), and F-score (Fu) metrics, as proposed by [5]. These metrics are designed to penalize false positives, with the F-score providing an overall performance score by combining precision and recall.

B. Results

Tables I and II present quantitative comparisons of our method with the state-of-the-art unsupervised motion segmentation method (EM) [21] on DAVIS-Moving and YTVOS-Moving datasets. The comparisons are limited to binary motion segmentation due to the availability of model weights

from [21]. However, these results are indicative of performance on multi-label segmentation tasks. Our method demonstrates superior performance on both datasets, producing higher F-scores than EM, showing more accurate segmentation results.

Figure 3 shows a qualitative comparison between our method and EM on the DAVIS-Moving and YTVOS-Moving datasets. The first row shows the original video frames, the second row shows the results from EM, the third row shows the results from our method, and the last row shows the ground truth. Both methods can detect moving objects in the scene, but our method produces much more coherent object masks and produces significantly less false-positive segments.

Method	Training Method	Pu	Ru	Fu
EM [21]	Unsupervised	61.29	86.56	68.90
Ours	Zero-Shot (no training)	74.47	77.78	75.96

TABLE I: Quantitative binary motion segmentation results of our method and the state-of-the-art unsupervised motion segmentation method (EM) [21] on DAVIS-Moving. Our method significantly outperforms EM with a much higher Fu score.

Method	Training Method	Pu	Ru	Fu
EM [21]	Unsupervised	41.78	39.06	35.38
Ours	Zero-Shot (no training)	54.63	50.36	50.78

TABLE II: Quantitative binary motion segmentation results of our method and EM [21] on the YTVOS-Moving dataset. Our method still significantly outperforms EM.

Table III shows quantitative comparison of our method and state-of-the-art supervised and semi-supervised methods

Method	Training Method	DAVIS-Moving			YTVOS-Moving		
		Pu	Ru	Fu	Pu	Ru	Fu
MoSeg [5]	Supervised	78.30	78.80	78.10	74.50	66.40	66.38
Raptor [7]	Supervised Features	75.90	79.67	75.93	64.43	60.94	60.35
RigidMask [6]		59.03	49.89	50.01	29.88	17.88	18.70
Ours	Zero-Shot (no training)	71.53	75.66	73.18	63.54	58.94	56.06

TABLE III: Quantitative results comparing our method with state-of-the-art supervised and semi-supervised methods on the DAVIS-Moving and YTVOS-Moving validation datasets show that while our method lags behind the supervised MoSeg, it matches semi-supervised methods and even outperforms RigidMask, despite no training.




Fig. 4: Qualitative comparison with state-of-the-art methods on DAVIS-Moving (rows 1-2) and YTVOS-Moving (rows 3-4). MoSeg performs best as a supervised method. RigidMask struggles with non-rigid motions, and Raptor has similar issues but to a lesser extent. Our method matches the performance of the supervised method in these challenging scenarios.

on the two benchmarks. Both binary and multi-label motion segmentation scenes are used. Although we did not conduct any training, our method achieves competitive results and even outperforms one of the semi-supervised methods (RigidMask).

We also show qualitative comparison with these methods in Figure 4. MoSeg [5] produces the best results on both datasets. RigidMask [6] fails to produce coherent segmentation on most scenes, and also fails to detect the motions of the parrot and the train on the last two rows, whereas our method successfully detects and segments them coherently. Raptor [7] is able to detect most objects in the scene thanks to its powerful semantic backbone, but it still over-segments non-rigid objects like the parrot. Our method performs well in these cases, performing almost as well as the supervised method.

Our method has two main limitations: slow inference speed due to multiple deep learning models, and the need for a predefined number of motions in the scene. The latter can

be addressed by incorporating an automatic model selection technique [59]. The inference speed, however, can only be improved by training an end-to-end network using our proposed motion residual functions.

C. Ablation Study

We present an ablation study to compare the performances of the two different motion models with the baseline, which is obtained from the raw object proposal alone. Both qualitative (Figure. 5) and quantitative (Table IV) comparisons are shown between our motion segmentation result obtained by only using optical flow as the motion cue vs. using both optical flow and the depth map, as well as the baseline results.

To produce the motion segmentation results only from optical flow, We use the optical flow motion model of EM [21], which is the state-of-the-art unsupervised motion segmentation method using only optical flow as input. Their 12-parameter quadratic parametric motion model integrates the unknown

Method	DAVIS-Moving			YTVOS-Moving		
	Pu	Ru	Fu	Pu	Ru	Fu
OC + Depth	71.53	75.66	73.18	63.54	58.94	56.06
OC	58.25	59.22	57.08	61.79	54.64	53.74
Base (obj. proposal)	43.17	86.24	52.12	48.49	73.01	50.82

TABLE IV: Quantitative ablation study: Motion segmentation results from only using optical flow (OC) as motion cue vs. using both optical flow and depth map (OC + Depth) as motion cue. Baseline results (Base) are also shown. Baseline is obtained by directly using the raw object proposals as the final motion segmentation mask. Bold numbers are the best results.




Fig. 5: Qualitative ablation study: Qualitative comparison between motion segmentation results using optical flow alone (OC) and both optical flow and depth map (OC + Depth). Pure optical flow based motion model (OC) suffers when multiple objects are at different depths. Combining optical flow with depth (OC + Depth) significantly mitigates this problem.

depth information as unknown parameters (as shown in equation (2)). Results indicate that the motion model integrating both optical flow and depth (OC + Depth) significantly outperforms the model relying solely on optical flow (OC) across all metrics on DAVIS-Moving. However, on the YTVOS-Moving dataset, the performance improvement is small, suggesting that unknown depth information is not a critical limiting factor. This can be attributed to several reasons: First, many objects labeled as moving in YTVOS-Moving are static in most frames. Second, the dataset includes significant occlusions and uncommon objects such as camouflaged animals. These challenges likely have a greater impact on the accuracy of motion segmentation than the absence of depth information.

V. CONCLUSION

We introduce a novel approach for dense monocular motion segmentation that operates without requiring any training. By integrating the deep learning models with traditional optical

flow-based methods, we propose a zero-shot technique that effectively clusters object proposals into distinct motion groups. Our method enhances the performance of conventional optical flow-based techniques by incorporating monocular depth maps, resulting in superior outcomes compared to using optical flow alone. Remarkably, despite the absence of training, our approach surpasses the state-of-the-art unsupervised motion segmentation methods on two widely-adopted benchmarks and rivals the top supervised and semi-supervised methods.

Future work will focus on enhancing our method by integrating additional motion cues and geometric models, such as keypoint correspondences and the fundamental matrix, to further boost its performance, as well as incorporating a model selection method to automatically infer the number of motions in the scene.

REFERENCES

- [1] R. I. Hartley and A. Zisserman, *Multiple View Geometry in Computer Vision*, 2nd ed. Cambridge University Press, ISBN: 0521540518, 2004.
- [2] E. Mohamed, M. Ewaisha, M. Siam, H. Rashed, S. Yogamani, W. Hamdy, M. El-Dakdouky, and A. El-Sallab, "Monocular Instance Motion Segmentation for Autonomous Driving: KITTI InstanceMotSeg Dataset and Multi-Task Baseline," in *2021 IEEE Intelligent Vehicles Symposium (IV)*. Nagoya, Japan: IEEE Press, Jul. 2021, pp. 114–121. [Online]. Available: <https://doi.org/10.1109/IV48863.2021.9575445>
- [3] J. Vertens, A. Valada, and W. Burgard, "SMSnet: Semantic motion segmentation using deep convolutional neural networks," in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Sep. 2017, pp. 582–589, iSSN: 2153-0866.
- [4] M. Ramzy, H. Rashed, A. E. Sallab, and S. Yogamani, "RST-MODNet: Real-time Spatio-temporal Moving Object Detection for Autonomous Driving," Dec. 2019, arXiv:1912.00438 [cs, stat] version: 1. [Online]. Available: <http://arxiv.org/abs/1912.00438>
- [5] A. Dave, P. Tokmakov, and D. Ramanan, "Towards Segmenting Anything That Moves," in *2019 IEEE/CVF International Conference on Computer Vision Workshop (ICCVW)*. Seoul, Korea (South): IEEE, Oct. 2019, pp. 1493–1502. [Online]. Available: <https://ieeexplore.ieee.org/document/9022103/>
- [6] G. Yang and D. Ramanan, "Learning to Segment Rigid Motions from Two Frames," in *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. Nashville, TN, USA: IEEE, Jun. 2021, pp. 1266–1275. [Online]. Available: <https://ieeexplore.ieee.org/document/9578593/>
- [7] M. Neoral, "Monocular Arbitrary Moving Object Discovery and Segmentation."
- [8] H. Sekkati and A. Mitiche, "A variational method for the recovery of dense 3D structure from motion," *Robotics and Autonomous Systems*, vol. 55, no. 7, pp. 597–607, Jul. 2007. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0921889006001941>
- [9] A. Wedel, A. Meißner, C. Rabe, U. Franke, and D. Cremers, "Detection and Segmentation of Independently Moving Objects from Dense Scene Flow," in *Energy Minimization Methods in Computer Vision and Pattern Recognition*, D. Cremers, Y. Boykov, A. Blake, and F. R. Schmidt, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2009, vol. 5681, pp. 14–27, series Title: Lecture Notes in Computer Science. [Online]. Available: http://link.springer.com/10.1007/978-3-642-03641-5_2
- [10] A. Papazoglou and V. Ferrari, "V.: Fast object segmentation in unconstrained video," in *In: ICCV (2013)*, 2013.
- [11] K. Fragkiadaki, Geng Zhang, and Jianbo Shi, "Video segmentation by tracing discontinuities in a trajectory embedding," in *2012 IEEE Conference on Computer Vision and Pattern Recognition*. Providence, RI: IEEE, Jun. 2012, pp. 1846–1853. [Online]. Available: <http://ieeexplore.ieee.org/document/6247883/>
- [12] M. Keuper, B. Andres, and T. Brox, "Motion Trajectory Segmentation via Minimum Cost Multicuts," in *2015 IEEE International Conference on Computer Vision (ICCV)*, Dec. 2015, pp. 3271–3279, iSSN: 2380-7504.
- [13] P. Bideau and E. Learned-Miller, "It's Moving! A Probabilistic Model for Causal Motion Segmentation in Moving Camera Videos," in *Computer Vision – ECCV 2016*, B. Leibe, J. Matas, N. Sebe, and M. Welling, Eds. Cham: Springer International Publishing, 2016, vol. 9912, pp. 433–449, series Title: Lecture Notes in Computer Science. [Online]. Available: http://link.springer.com/10.1007/978-3-319-46484-8_26
- [14] A. Mitiche and J. Aggarwal, *Computer Vision Analysis of Image Motion by Variational Methods*, ser. Springer Topics in Signal Processing. Cham: Springer International Publishing, 2014, vol. 10. [Online]. Available: <http://link.springer.com/10.1007/978-3-319-00711-3>
- [15] S. Negahdaripour and B. K. P. Horn, "Direct Passive Navigation," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. PAMI-9, no. 1, pp. 168–176, Jan. 1987. [Online]. Available: <https://ieeexplore.ieee.org/document/4767884>
- [16] P. Bideau, A. RoyChowdhury, R. R. Menon, and E. Learned-Miller, "The Best of Both Worlds: Combining CNNs and Geometric Constraints for Hierarchical Motion Segmentation," in *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*. Salt Lake City, UT, USA: IEEE, Jun. 2018, pp. 508–517. [Online]. Available: <https://ieeexplore.ieee.org/document/8578158/>
- [17] M. Siam, H. Mahgoub, M. Zahran, S. Yogamani, M. Jagersand, and A. El-Sallab, "MODNet: Motion and Appearance based Moving Object Detection Network for Autonomous Driving," in *2018 21st International Conference on Intelligent Transportation Systems (ITSC)*, Nov. 2018, pp. 2859–2864, iSSN: 2153-0017.
- [18] M. Bosch, "Deep Learning for Robust Motion Segmentation with Non-Static Cameras," Feb. 2021, arXiv:2102.10929 [cs]. [Online]. Available: <http://arxiv.org/abs/2102.10929>
- [19] P. Bideau, R. R. Menon, and E. Learned-Miller, "MoA-Net: Self-supervised Motion Segmentation," in *Computer Vision – ECCV 2018 Workshops*, L. Leal-Taixé and S. Roth, Eds. Cham: Springer International Publishing, 2019, vol. 11134, pp. 715–730, series Title: Lecture Notes in Computer Science. [Online]. Available: http://link.springer.com/10.1007/978-3-030-11024-6_55
- [20] M. Faisal, I. Akhter, M. Ali, and R. Hartley, "EpO-Net: Exploiting Geometric Constraints on Dense Trajectories for Motion Saliency," in *2020 IEEE Winter Conference on Applications of Computer Vision (WACV)*. Snowmass Village, CO, USA: IEEE, Mar. 2020, pp. 1873–1882. [Online]. Available: <https://ieeexplore.ieee.org/document/9093589/>
- [21] E. Meunier, A. Badoual, and P. Bouthemy, "EM-Driven Unsupervised Learning for Efficient Motion Segmentation," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 45, no. 4, pp. 4462–4473, Apr. 2023, conference Name: IEEE Transactions on Pattern Analysis and Machine Intelligence.
- [22] A. Delong, A. Osokin, H. N. Isack, and Y. Boykov, "Fast approximate energy minimization with label costs," in *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, Jun. 2010, pp. 2173–2180, iSSN: 1063-6919.
- [23] H. Isack and Y. Boykov, "Energy-Based Geometric Multi-model Fitting," *International Journal of Computer Vision*, vol. 97, no. 2, pp. 123–147, Apr. 2012. [Online]. Available: <http://link.springer.com/10.1007/s11263-011-0474-7>
- [24] D. Hutchison, T. Kanade, J. Kittler, J. M. Kleinberg, F. Mattern, J. C. Mitchell, M. Naor, O. Nierstrasz, C. Pandu Rangan, B. Steffen, M. Sudan, D. Terzopoulos, D. Tygar, M. Y. Vardi, G. Weikum, T. Brox, and J. Malik, "Object Segmentation by Long Term Analysis of Point Trajectories," in *Computer Vision – ECCV 2010*, K. Daniilidis, P. Maragos, and N. Paragios, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2010, vol. 6315, pp. 282–295, series Title: Lecture Notes in Computer Science. [Online]. Available: http://link.springer.com/10.1007/978-3-642-15555-0_21
- [25] T. Brox and J. Malik, "Object segmentation by long term analysis of point trajectories," in *Proceedings of the 11th European conference on Computer vision: Part V*, ser. ECCV'10. Berlin, Heidelberg: Springer-Verlag, Sep. 2010, pp. 282–295.
- [26] E. Elhamifar and R. Vidal, "Sparse Subspace Clustering: Algorithm, Theory, and Applications," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 35, no. 11, pp. 2765–2781, Nov. 2013. [Online]. Available: <http://ieeexplore.ieee.org/document/6482137/>
- [27] P. Ochs, J. Malik, and T. Brox, "Segmentation of Moving Objects by Long Term Video Analysis," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 36, no. 6, pp. 1187–1200, Jun. 2014, conference Name: IEEE Transactions on Pattern Analysis and Machine Intelligence.
- [28] T. Lai, H. Wang, Y. Yan, T.-J. Chin, and W.-L. Zhao, "Motion Segmentation Via a Sparsity Constraint," *IEEE Transactions on Intelligent Transportation Systems*, vol. 18, no. 4, pp. 973–983, Apr. 2017, conference Name: IEEE Transactions on Intelligent Transportation Systems.
- [29] X. Xu, L. F. Cheong, and Z. Li, "Motion Segmentation by Exploiting Complementary Geometric Models," in *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*. Salt Lake City, UT, USA: IEEE, Jun. 2018, pp. 2859–2867. [Online]. Available: <https://ieeexplore.ieee.org/document/8578400/>
- [30] F. Arrigoni, L. Magri, and T. Pajdla, "On the Usage of the Trifocal Tensor in Motion Segmentation," in *Computer Vision – ECCV 2020*, A. Vedaldi, H. Bischof, T. Brox, and J.-M. Frahm, Eds. Cham: Springer International Publishing, 2020, vol. 12365, pp. 514–530, series Title: Lecture Notes in Computer Science. [Online]. Available: https://link.springer.com/10.1007/978-3-030-58565-5_31
- [31] R. Yao, G. Lin, S. Xia, J. Zhao, and Y. Zhou, "Video Object Segmentation and Tracking: A Survey," *ACM Transactions on Intelligent Systems and Technology*, vol. 11, no. 4, pp. 36:1–36:47, May 2020. [Online]. Available: <https://dl.acm.org/doi/10.1145/3391743>

- [32] Aloimonos and C. M. Brown, "Direct processing of curvilinear sensor motion from a sequence of perspective images," 1984, volume: 72.
- [33] B. K. P. Horn and E. J. Weldon, "Direct methods for recovering motion," *International Journal of Computer Vision*, vol. 2, no. 1, pp. 51–76, Jun. 1988. [Online]. Available: <https://doi.org/10.1007/BF00836281>
- [34] R. Vidal and D. Singaraju, "A closed form solution to direct motion segmentation," in *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05)*, vol. 2, Jun. 2005, pp. 510–515 vol. 2, iSSN: 1063-6919.
- [35] Z. Teed and J. Deng, "RAFT: Recurrent All-Pairs Field Transforms for Optical Flow," in *Computer Vision – ECCV 2020*, A. Vedaldi, H. Bischof, T. Brox, and J.-M. Frahm, Eds. Cham: Springer International Publishing, 2020, vol. 12347, pp. 402–419, series Title: Lecture Notes in Computer Science. [Online]. Available: https://link.springer.com/10.1007/978-3-030-58536-5_24
- [36] D. Sun, C. Herrmann, F. Reda, M. Rubinstein, D. J. Fleet, and W. T. Freeman, "Disentangling Architecture and Training for Optical Flow," in *Computer Vision – ECCV 2022*, S. Avidan, G. Brostow, M. Cissé, G. M. Farinella, and T. Hassner, Eds. Cham: Springer Nature Switzerland, 2022, vol. 13682, pp. 165–182, series Title: Lecture Notes in Computer Science. [Online]. Available: https://link.springer.com/10.1007/978-3-031-20047-2_10
- [37] S. Choudhury, L. Karazija, I. Laina, A. Vedaldi, and C. Rupprecht, "Guess What Moves: Unsupervised Video and Image Segmentation by Anticipating Motion," May 2022, arXiv:2205.07844 [cs]. [Online]. Available: <http://arxiv.org/abs/2205.07844>
- [38] E. Meunier, A. Badoual, and P. Boutheymy, "EM-driven unsupervised learning for efficient motion segmentation," Mar. 2022, arXiv:2201.02074 [cs]. [Online]. Available: <http://arxiv.org/abs/2201.02074>
- [39] P. H. S. Torr, "Geometric motion segmentation and model selection," *Philosophical Transactions of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences*, vol. 356, no. 1740, pp. 1321–1340, May 1998. [Online]. Available: <https://royalsocietypublishing.org/doi/10.1098/rsta.1998.0224>
- [40] D. Barath and J. Matas, "Progressive-X: Efficient, Anytime, Multi-Model Fitting Algorithm," in *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*. Seoul, Korea (South): IEEE, Oct. 2019, pp. 3779–3787. [Online]. Available: <https://ieeexplore.ieee.org/document/9010674/>
- [41] Y. Jiang, Q. Xu, K. Ma, Z. Yang, X. Cao, and Q. Huang, "What to Select: Pursuing Consistent Motion Segmentation from Multiple Geometric Models," *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 35, no. 2, pp. 1708–1716, May 2021, number: 2. [Online]. Available: <https://ojs.aaai.org/index.php/AAAI/article/view/16264>
- [42] Z. Xi, J. Liu, B. Luo, and Q. Qin, "Multi-Motion Segmentation: Combining Geometric Model-Fitting and Optical Flow for RGB Sensors," *IEEE Sensors Journal*, vol. 22, no. 7, pp. 6952–6963, Apr. 2022, conference Name: IEEE Sensors Journal.
- [43] S. Rao, R. Tron, R. Vidal, and Y. Ma, "Motion Segmentation in the Presence of Outlying, Incomplete, or Corrupted Trajectories," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 32, no. 10, pp. 1832–1845, Oct. 2010, conference Name: IEEE Transactions on Pattern Analysis and Machine Intelligence.
- [44] R. Tron and R. Vidal, "A Benchmark for the Comparison of 3-D Motion Segmentation Algorithms," in *2007 IEEE Conference on Computer Vision and Pattern Recognition*, Jun. 2007, pp. 1–8, iSSN: 1063-6919. [Online]. Available: <https://ieeexplore.ieee.org/abstract/document/4269999>
- [45] R. Vidal, "Subspace Clustering," *IEEE Signal Processing Magazine*, vol. 28, no. 2, pp. 52–68, Mar. 2011. [Online]. Available: <http://ieeexplore.ieee.org/document/5714408/>
- [46] Y. Zhang, X. Huang, J. Ma, Z. Li, Z. Luo, Y. Xie, Y. Qin, T. Luo, Y. Li, S. Liu, Y. Guo, and L. Zhang, "Recognize Anything: A Strong Image Tagging Model," Jun. 2023, arXiv:2306.03514 [cs]. [Online]. Available: <http://arxiv.org/abs/2306.03514>
- [47] S. Liu, Z. Zeng, T. Ren, F. Li, H. Zhang, J. Yang, C. Li, J. Yang, H. Su, J. Zhu, and L. Zhang, "Grounding DINO: Marrying DINO with Grounded Pre-Training for Open-Set Object Detection," Mar. 2023, arXiv:2303.05499 [cs]. [Online]. Available: <http://arxiv.org/abs/2303.05499>
- [48] F. Rajič, L. Ke, Y.-W. Tai, C.-K. Tang, M. Danelljan, and F. Yu, "Segment Anything Meets Point Tracking," Jul. 2023, arXiv:2307.01197 [cs]. [Online]. Available: <http://arxiv.org/abs/2307.01197>
- [49] Z. Yang and Y. Yang, "Decoupling Features in Hierarchical Propagation for Video Object Segmentation," *Advances in Neural Information Processing Systems*, vol. 35, pp. 36 324–36 336, Dec. 2022. [Online]. Available: https://proceedings.neurips.cc/paper_files/paper/2022/hash/eb890c36af87e4ca82e8ef7bcb6a284-Abstract-Conference.html
- [50] Y. Cheng, L. Li, Y. Xu, X. Li, Z. Yang, W. Wang, and Y. Yang, "Segment and Track Anything," May 2023, arXiv:2305.06558 [cs]. [Online]. Available: <http://arxiv.org/abs/2305.06558>
- [51] D. Sun, C. Herrmann, F. Reda, M. Rubinstein, D. J. Fleet, and W. T. Freeman, "Disentangling Architecture and Training for Optical Flow," in *Computer Vision – ECCV 2022*, ser. Lecture Notes in Computer Science, S. Avidan, G. Brostow, M. Cissé, G. M. Farinella, and T. Hassner, Eds. Cham: Springer Nature Switzerland, 2022, pp. 165–182.
- [52] M. Oquab, T. Darcet, T. Moutakanni, H. Vo, M. Szafraniec, V. Khalidov, P. Fernandez, D. Haziza, F. Massa, A. El-Nouby, M. Assran, N. Ballas, W. Galuba, R. Howes, P.-Y. Huang, S.-W. Li, I. Misra, M. Rabbat, V. Sharma, G. Synnaeve, H. Xu, H. Jegou, J. Mairal, P. Labatut, A. Joulin, and P. Bojanowski, "DINOv2: Learning Robust Visual Features without Supervision," Apr. 2023. [Online]. Available: <https://arxiv.org/abs/2304.07193v1>
- [53] H. C. Longuet-Higgins and K. Prazdny, "The Interpretation of a Moving Retinal Image," *Proceedings of the Royal Society of London. Series B, Biological Sciences*, vol. 208, no. 1173, pp. 385–397, 1980, publisher: The Royal Society. [Online]. Available: <https://www.jstor.org/stable/35316>
- [54] T.-j. Chin, H. Wang, and D. Suter, "The Ordered Residual Kernel for Robust Motion Subspace Clustering," in *Advances in Neural Information Processing Systems*, vol. 22. Curran Associates, Inc., 2009. [Online]. Available: https://proceedings.neurips.cc/paper_files/paper/2009/hash/b337e84de8752b27eda3a12363109e80-Abstract.html
- [55] U. Von Luxburg, "A tutorial on spectral clustering," *Statistics and Computing*, vol. 17, no. 4, pp. 395–416, Dec. 2007. [Online]. Available: <http://link.springer.com/10.1007/s11222-007-9033-z>
- [56] Y. Huang and J. Zelek, "Motion Segmentation from a Moving Monocular Camera," Sep. 2023, arXiv:2309.13772 [cs]. [Online]. Available: <http://arxiv.org/abs/2309.13772>
- [57] J. Pont-Tuset, F. Perazzi, S. Caelles, P. Arbeláez, A. Sorkine-Hornung, and L. Van Gool, "The 2017 DAVIS Challenge on Video Object Segmentation," Mar. 2018, arXiv:1704.00675 [cs]. [Online]. Available: <http://arxiv.org/abs/1704.00675>
- [58] N. Xu, L. Yang, Y. Fan, J. Yang, D. Yue, Y. Liang, B. Price, S. Cohen, and T. Huang, "YouTube-VOS: Sequence-to-Sequence Video Object Segmentation," in *Computer Vision – ECCV 2018*, V. Ferrari, M. Hebert, C. Sminchisescu, and Y. Weiss, Eds. Cham: Springer International Publishing, 2018, vol. 11209, pp. 603–619, series Title: Lecture Notes in Computer Science. [Online]. Available: https://link.springer.com/10.1007/978-3-030-01228-1_36
- [59] Y. Huang and J. Zelek, "A Unified Model Selection Technique for Spectral Clustering Based Motion Segmentation," *Journal of Computational Vision and Imaging Systems*, vol. 9, no. 1, pp. 68–71, 2023, number: 1. [Online]. Available: <https://openjournals.uwaterloo.ca/index.php/vsl/article/view/5870>