A Hybrid Sparse-Dense Monocular SLAM System for Autonomous Driving

Louis Gallagher¹, Varun Ravi Kumar², Senthil Yogamani³ and John B. McDonald¹

Abstract-In this paper, we present a system for incrementally reconstructing a dense 3D model of the geometry of an outdoor environment using a single monocular camera attached to a moving vehicle. Dense models provide a rich representation of the environment facilitating higher-level scene understanding, perception, and planning. Our system employs dense depth prediction with a hybrid mapping architecture combining state-of-the-art sparse features and dense fusionbased visual SLAM algorithms within an integrated framework. Our novel contributions include design of hybrid sparsedense camera tracking and loop closure, and scale estimation improvements in dense depth prediction. We use the motion estimates from the sparse method to overcome the large and variable inter-frame displacement typical of outdoor vehicle scenarios. Our system then registers the live image with the dense model using whole-image alignment. This enables the fusion of the live frame and dense depth prediction into the model. Global consistency and alignment between the sparse and dense models are achieved by applying pose constraints from the sparse method directly within the deformation of the dense model. We provide qualitative and quantitative results for both trajectory estimation and surface reconstruction accuracy, demonstrating competitive performance on the KITTI dataset. Qualitative results of the proposed approach are illustrated in https://youtu.be/Pn2uaVqjskY. Source code for the project is publicly available at the following repository https: //github.com/robotvisionmu/DenseMonoSLAM

I. INTRODUCTION

Over the past decade, approaches to fusion-based dense visual SLAM have demonstrated the ability to build high fidelity dense 3D models of an environment facilitating higher-level scene understanding, perception and planning [1], [2], [3], [4]. However, given that these techniques typically require active RGB-D sensors, their applicability in outdoor scenarios has been limited. On the other hand, sparse and semi-dense monocular systems have found a large degree of success in these environments but are limited in their resulting sparse or partially dense representations [5], [5], [6], [7], [8], [9], [10]. Here, we investigate the potential of recent results in monocular dense depth prediction in

*Research presented in this paper was supported by IRC grant GOIPG/2016/1320 HEA COVID extension fund, Science Foundation Ireland grant 13/RC/2094 to Lero - the Irish Software Research Centre and Science Foundation Ireland grant 16/RI/3399.

¹Louis Gallagher and John McDonald are with Lero - the Irish Software Research Centre and the Department of Computer Science, Maynooth University, Maynooth, Ireland {Louis.Gallagher, John.McDonald}@mu.ie

²Varun Ravi Kumar is with Valeo DAR Kronach, Germany. varun.ravi-kumar@valeo.com

³Senthil Yogamani is with Valeo Vision Systems, Ireland senthil.yogamani@valeo.com 978-1-6654-1213-1/21/\$31.00 ©2021 IEEE



Fig. 1: Trajectory and dense reconstruction of an odometry sequence from [11], [12]. (a) shows the final surfel model produced by our system. The model contains approximately 16m surfels. (b) shows the estimated trajectory alongside the ground truth trajectory. (c) shows an up-close view of the reconstruction as the vehicle passes through an intersection contained in (a).

combining these approaches' strengths. In particular, we employ dense depth prediction in developing a monocular SLAM system that achieves comparable accuracy to sparse camera tracking algorithms while allowing dense fusionbased modeling of the environment. We demonstrate the effectiveness of this approach by tracking and mapping an environment from a single monocular camera attached to a moving vehicle.

Using passive vision sensors over active sensors, such as LiDAR, in autonomous driving SLAM systems has many advantages. A calibrated monocular camera is inexpensive, lightweight, and can be used as a direction sensor providing rich photometric and geometric measurements of the scene at every frame. In contrast, LiDAR is expensive, heavy, and produces a relatively sparse signal compared to the dense measurements provided by a vision sensor. Indeed, recent deep learning-based depth estimation and object detection indicates that the accuracy advantage of LiDAR over monocular vision-based systems is closing [13].

On the other hand, monocular vision systems suffer from ambiguity and drift in the estimated scale of the reconstruction [14] and are not as robust as stereo and multi-sensor systems in texture-less regions and during fast camera motion [6], [7]. Furthermore, with the exception of a few techniques for dense and semi-dense map representations [15], [8], [16], these systems produce sparse reconstructions which only offer a coarse-grained description of the scene's geometry [17], [18], [19], [5], [20].

Multi-sensor platforms attempt to leverage the combined advantages of monocular, inertial, stereo, LiDAR, and GPS sensors. However, such platforms need careful internal and external calibration, which can be done online to maintain synchronization and geometric alignment between the sensors and, ultimately, achieve accurate performance. In the autonomous driving domain, even the most carefully calibrated rig will move and warp in real-time due to the effects of heat, wear and tear of mechanical parts like tires, variations in loading of the vehicle, and the knocks and bumps inherent in driving on roads [21]. GPS comes with its complexities, requiring line-of-sight with at least four satellites to infer position, which can prove challenging in urban environments [22].

More recently, several systems have emerged that represent the state-of-the-art in traditional monocular SLAM, and visual-odometry performance [6], [20], [23], [24], [9], [10], [8]. The ORB-SLAM open-source SLAM frameworks [5], [20], [6] in particular represent the state-of-the-art in monocular SLAM. Similar to PTAM, ORB-SLAM [20] splits the SLAM problem into sub-problems that can be solved in different threads; (i) a feature-based camera tracking thread that runs in real-time; (ii) a local mapping thread that takes keyframes extracted by the first thread and optimizes a local area of the map around this keyframe using bundle-adjustment (BA); and (iii) a global mapping thread that integrates loop closure constraints and maintains global consistency.

Visual SLAM systems based around the concepts of dense alternation, every pixel fusion, and direct wholeimage alignment-based camera tracking, have shown that it is possible to build dense, high-fidelity models of indoor scenes in real-time with active depth sensors such as the Microsoft Kinect. These systems and the models they produce have many benefits in downstream machine perceptions and spatial reasoning tasks. They can produce 'watertight' maps that model the continuous surface that underlies the scene. For example, this allows them to synthesize accurate novel views and realistically fuse augmentations into the scene for AR effects [1], [3], [2], [4], [25]. High-resolution dense models may become more necessary for robot planning in the future to accommodate a broader range of scenarios (e.g., analyzing surface geometry in unevenly paved roads [26], obstacle detection and avoidance [13], etc.).

More and more deep learning is being used to improve the performance of SLAM systems. Broadly speaking, neural networks have been applied in two critical areas within the SLAM problem; in the representation of geometric and visual information [27], [28], [29], and in overcoming the ill-posedness of many of SLAM's subproblems, for example, depth estimation from monocular images [30], [31], [32], [33]. Several systems have been presented that demonstrate hybrid approaches to monocular visual SLAM and VO [34], [27], [35], [36], while other approaches implement a fully learned SLAM pipeline [37], [38]. In comparison to this work, these systems have either not been designed with large-scale fast, outdoor camera motion in mind [34], [27], [37], [35], or do not produce dense models [36], [38].

The work most similar to our aproach is the dense surfelbased fusion system of [39], combining camera tracking from a sparse SLAM system (ORB-SLAM2 [20]) with a dense fusion-based back-end. The substantial difference with the current work is the degree to which they couple the sparse and dense maps. In [39], a sub-mapping approach is employed where each surfel is anchored to a reference keyframe from the sparse system. Local and global optimisations of the keyframe poses by the sparse system are reflected in the dense map by ensuring the pose of a surfel relative to its anchor keyframe remains constant. In contrast, our system maintains a much looser coupling between the sparse and dense representations. The sparse pose is used to initialise a dense alignment that registers the camera to the dense map allowing for fusion of the current frame. Loop closures from the sparse system are used to constrain a deformation graph that non-rigidly deforms the surfels themselves. In this way, the benefits of each system are leveraged, while the systems themselves are only indirectly connected through 6-DOF transformations and 3D constraints. Furthermore, we quantitatively demonstrate the effectiveness of our approach in outdoor automotive environments using deep-learned depth estimation.

The contribution of this paper is the development of a system for live dense metric 3D reconstruction in outdoor environments using a single RGB camera attached to a moving vehicle. Although other researchers have reported dense and semi-dense SLAM or dense VO approaches, to the best of our knowledge, this is the first fully dense SLAM system quantitatively evaluated on automotive scenarios. The novelty of our approach lies in the use of a dense depth prediction network [30] within a hybrid architecture, combining state-of-the-art sparse feature tracking and dense fusion-based visual mapping algorithms in a loosely coupled fashion. We improve our previous approach to dense depth prediction [30] for the SLAM application by the addition of new regularization losses and better scale estimation. Using the motion estimates from the sparse method to overcome the large and variable inter-frame displacement typical of outdoor vehicle scenarios, our system then registers the live image with the dense model using whole-image alignment. This allows dense depth prediction and fusion of the live frame into the model. Global consistency and alignment between the sparse and dense models are achieved by applying pose constraints from the sparse method directly within the deformation of the dense model. We evaluate the method over the KITTI benchmark dataset, providing qualitative and quantitative results for both the trajectory and surface reconstruction accuracy.

II. A HYBRID APPROACH

We take a hybrid approach to dense monocular tracking and mapping. A feature-based SLAM system, ORB-SLAM3[6], is used to provide an initial estimate of the camera's pose for each frame. The pipeline then follows a dense alternation architecture, extending the ElasticFusion (EF) SLAM system [3], in which the map is first held fixed. At the same time, the camera is tracked against it, using the initial pose estimate from the previous step. Once the camera pose has been estimated, it can fuse the current frame into the map. We densely predict per pixel metric depth estimates using a SOTA self-supervised convolutional neural network, UnRectDepthNet [30]. The various subsystems predominantly operate on different processors; the GPU is used extensively by the dense alternation and depth prediction network, while ORB-SLAM operates on the CPU. Our hybrid architecture is summarised as follows:

- A scale-aware depth prediction network is used to estimate a metric depth map for each frame. An initial estimate of the camera motion is made using ORB-SLAM's feature-based camera tracker, which is suited to the fast motion of the vehicle.
- 2) The initial pose estimate is further refined by aligning it to the current active model in view of the camera.
- 3) Live RGB images and corresponding predicted metric depth maps are fused into a global dense surfel model of the scene. The surfel model is divided into an active and an inactive portion as per the original EF algorithm.
- 4) When ORB-SLAM identifies a loop closure, we use the resulting loop closure constraint within the EF deformation graph to correct the geometry of the dense surfaces. This brings the previously visited inactive portion of the map back into alignment with the current active portion. Importantly, this also keeps the different map and camera trajectories of ORB-SLAM and EF consistent.

Figure 2 shows the architecture of our system. In the remainder of this section, we will describe how these key elements are combined into one system that builds consistent dense surfel models from a monocular camera attached to a moving vehicle. The section provides complete coverage of the system where the principal novelty is covered in II-A, II-D, and II-E. We then present and discuss experimental results that demonstrate the efficacy of our system in III.

A. Scale-aware depth estimation

Following UnRectDepthNet [30], we establish the same *structure-from-motion (SfM)* framework for self-supervised depth estimation. We carry out the view synthesis by employing the pinhole camera projection model. The final objective consists of a photometric term \mathcal{L}_p and an edge smoothness regularization term \mathcal{L}_s . In addition, the cross-sequence depth consistency loss \mathcal{L}_{dc} and the scale recovery approach are employed. In the following paragraphs, we discuss the new improvements that lead to a considerable gain in accuracy.

We employ feature-metric losses from [40], where discriminative, \mathcal{L}_{dis} , and convergent, \mathcal{L}_{cvt} , losses are calculated on the current frame's, I_t 's, feature representation by incorporating a self-attention autoencoder to obtain robust global features of the scene. The primary goal of \mathcal{L}_{dis} and \mathcal{L}_{cvt} is to keep the optimization objective from getting trapped at several local minima for low texture areas such as sky and road. It is a crucial loss feature that uses image gradients to penalize small slopes while emphasizing low-texture areas. The self-supervised loss landscapes are restricted from forming proper convergence basins using first-order derivatives to regularize the target features. However, enforcing discriminative loss alone will not ensure that we reach the best solution during gradient descent. As there exists inconsistency among first-order gradients, (i.e., gradients that are spatially adjacent point in opposite directions), convergent \mathcal{L}_{cvt} loss is employed to enable gradient descent from a faroff distance. It has a relatively large convergence radius and expresses the loss to have uniform gradients throughout the optimization step by supporting feature gradient smoothness and large convergence radii accordingly. The total objective \mathcal{L}_{depth} is

$$\mathcal{L}_{depth} = \mathcal{L}_r(I_t, \hat{I}_{t' \to t}) + \alpha \ \mathcal{L}_s(\hat{D}_t) + \rho \ \mathcal{L}_{dc}(\hat{D}_t, \hat{D}_{t'}) \quad (1) + \mathcal{L}_r(\hat{F}_t, \hat{F}_{t' \to t}) + \zeta \ \mathcal{L}_{dis}(I_t, \hat{F}_t) + \eta \ \mathcal{L}_{cvt}(I_t, \hat{F}_t)$$

1

Where \hat{F}_t is the estimated feature at time t, I_t is the current colour image, \hat{D}_t the estimated depth map, \mathcal{L}_r is the standard reconstruction matching term, and α , ρ , ζ and η weigh the smoothness term \mathcal{L}_s , cross-sequence depth consistency \mathcal{L}_{dc} , discriminative \mathcal{L}_{dis} and convergent \mathcal{L}_{cvt} losses respectively.

Scale ambiguity is a challenging problem in monocular depth estimation [41], [42]. As a result, an absolute value is needed to serve as an anchor point, provided through a measurement from another dedicated sensor to obtain actual depth estimation. We use a combination of Velodyne point cloud and calibration information to improve the estimation of scale. We use Velodyne Lidar as ground truth and estimate the scale factor by associating the calculation with the correct pixel in the image plane, minimizing ℓ_2 loss. In addition, we assume depth consistency across training and testing datasets. It can be helpful in datasets with high pose variability, such as KITTI, where the camera is still at the same height and looking at the ground from the same perspective. We fit a ground plane and use calibration information and known camera height to find the scale factor. We use a combination to improve our scale estimation.



Fig. 2: An overview of our system architecture. UnRectDepthNet is used to estimate a depth map for the current live camera image. ORB-SLAM then tracks the motion of the camera since the last frame. Internally ORB-SLAM continues to extract keyframes and pass them to its local and global mapping threads. The sparse feature map on the right is used internally by ORB-SLAM. Once the camera motion has been estimated, the pose is used to synthesize a dense view of the current active surfel model. The pose estimate is further refined by performing a direct joint photometric and geometric alignment between the synthetic and live RGB-D frames, taking the pose estimate from ORB-SLAM as an initialization point. Assuming no local or global loop closures have occurred during this timestep, the live RGB image and predicted depth are fused into the surfel model. If ORB-SLAM has detected a global loop closure, then a global deformation is attempted. Surface-to-surface constraints for optimization of the deformation graph are generated from the pre-corrected and corrected (post loop closure) poses of the camera P_{kf}^{t} and \hat{P}_{kf}^{t} respectively. If deformation graph optimization fails or no global loop is detected by ORB-SLAM, a local loop closure is attempted as per the original EF algorithm.

B. ORB-SLAM3 - Feature-based RGBD Tracking

ORB-SLAM3 [6] builds a sparse map of a scene, represented using a *covisibility* graph where each node in the graph corresponds to a keyframe consisting of a pose and a set of 3D points. When two keyframes view common map points, an edge between them is added to the graph. The system has 3 main threads. The camera tracking thread receives RGB-D frames from the camera, extracts ORB features, and computes an initial pose estimate via motion only BA with the previous frame. The estimate is further refined by aligning the current frame to a local map of covisible keyframes. New keyframes are detected and sent to a background thread. Here the local map in the vicinity of the new keyframe is optimized with a full bundle adjustment. Loops between the latest keyframe and historic keyframes are detected with DBoW[43]. If geometric alignment between the new keyframe and the matched keyframe succeeds, then a loop is closed, and an edge between the two keyframes is added to the graph. The local map in the vicinity of the keyframe is rigidly transformed into place. The rest of the keyframe graph is corrected using pose graph optimization. A final full BA is performed to recover the MAP estimate of all keyframe poses and structure.

C. ElasticFusion - Dense Surfel Fusion

ElasticFusion [3] estimates a dense surfel map \mathcal{M} of a scene as viewed from a handheld RGB-D sensor. The map

 \mathcal{M} is a flat list structure containing a set of surface elements \mathcal{M}^s describing local planar patches of the underlying scene surface. Each \mathcal{M}^s contains a 3D position, a normal, a radius r indicating its extent, a confidence value indicating quality of its estimation, an RGB color, a timestamp t when it was last fused, and a timestamp t_o when it was inserted into the map. \mathcal{M} is itself divided into two disjoint subsets; an active portion containing surfels that have recently been measured by the live camera data and an inactive portion containing those surfels that have not been measured in a threshold amount of time. Camera tracking is performed by direct alignment of the current camera frame to a synthetic model view using the pose of the camera at the previous timestep. Frames are fused by back projecting each pixel to a 3D point using the depth map and the camera's intrinsic matrix. The system also has a visual place recognition loop closure mechanism based on fern encoding; however, we do not utilize this module. Local loops are identified by aligning views of the active and inactive portions of \mathcal{M}^s in the current camera view, allowing inactive surfels to be reactivated. Both global and local loop closures provide surface-to-surface constraints to optimize a space deforming graph, which is then applied to \mathcal{M}^s . Camera tracking and deformation-based loop closures will be discussed in later sections as we describe how we combine them with the other systems.



Fig. 3: A loop closure example; (i) a car begins exploring (green arrow); (ii) after a period of time a loop is closed between an area previously mapped by the car and its current location; (iii) before a loop closure is applied, drift is evident in the model; (iv) the active portion of the model (green) has drifted from the inactive portion of the model (gray). Red arrows highlight corresponding points between the active and inactive map; (v) a global surface deformation has been applied bringing the active portion of the map back into alignment with the inactive portion. Red lines are surface-to-surface constraints generated from the loop closing pair $P_{kf}^t \rightarrow \hat{P}_{kf}^t$; (vi) as the car moves back into the region of the road already mapped in the first pass, local loop closures are triggered, reactivating portions of the map for use in mapping and tracking; (vii) the final global model.

D. Hybrid Camera Tracking

In our system, at each timestep, t, the current live camera frame \mathcal{F}_t , consisting of color image \mathbf{I}_t and predicted depth map \mathbf{D}_{t}^{p} , is passed to ORB-SLAM to compute an initial estimate of the camera pose \mathbf{P}_t . Note that this is the camera pose before any refinement by the local mapping thread in the case that \mathcal{F}_t is extracted as a keyframe. Though this initial pose is accurate (as can be seen in Table I), it cannot be used directly for the fusion of \mathcal{F}_t into \mathcal{M} . ORB-SLAM continues to improve its local map, while the dense surfel map does not undergo any significant correction until a local or global loop closure is triggered. This can lead to P_t , and therefore the live frame, becoming misaligned from the dense map. Our goal is to track the camera *and* to build an accurate model, and so we need to balance the accuracy of \mathbf{P}_t against the need to stay aligned with the dense model in order to fuse \mathcal{F}_t accurately.

To bring the camera back into alignment with the model, we perform a frame-to-model refinement of \mathbf{P}_t . The active map around \mathbf{P}_t is rendered into a virtual camera positioned at \mathbf{P}_t . A 6DOF transformation $T \in SE_3$ that aligns the live frame to the virtual one is then estimated in an iterative non-linear least-squares joint photometric and geometric optimization embedded in a 3 level image pyramid. Composing T with \mathbf{P}_t yields a refined pose estimate for the current frame that can now be used to fusion as per the original EF algorithm [3].

E. Hybrid Loop Closures

When ORB-SLAM identifies a loop closure between the latest keyframe Kf_i and a historic keyframe in the map, Kf_h , a transformation T_i^h that aligns Kf_i with Kf_h is computed. This yields a pose pair; the uncorrected pose of the new keyframe \mathbf{P}_{kf}^i and the pose of the keyframe after

loop correction $\hat{\mathbf{P}}_{kf}^{i}$. ORB-SLAM's loop closure mechanism proceeds to apply the loop closure and optimize its sparse map as per the original algorithm [6].

To update the dense surface to reflect the loop closure $\mathbf{P}_{kf}^i \to \hat{\mathbf{P}}_{kf}^i$, we non-rigidly deform each surfel using a deformation graph \mathcal{G} in the same way as EF [3]. \mathcal{G} consists of a set of nodes \mathcal{G}^n sampled from the surfels in \mathcal{M} . Each node consists of a position (i.e., the sampled surfel's position), a timestamp (also sampled surfel's timestamp), an affine transformation \mathcal{G}_T^n and a set of temporally nearby neighbor nodes. The temporal connectivity of the nodes prevents interference between multiple passes of the same surface at different points in time. To apply \mathcal{G} to \mathcal{M} , a set of temporally and spatially nearby influencing nodes is computed for each $\mathcal{M}^s \in \mathcal{M}$. A weighted sum of the \mathcal{G}_T^n in these nodes are applied to the surfel to deform it into place.

To optimize the \mathcal{G}_T^n , a set of surface-to-surface correspondences are generated from a loop closure, mapping surface points from the active region to the inactive region of the model. Cost function terms on the distance between correspondences, as well as terms that maximize the rigidity and smoothness of the deformation, ensure that, when applied to the surface, it brings the correspondences together. A pinning term holds the inactive map in place, ensuring that it is the active map that is deformed into the inactive map. Further constraints on the relative position of historical correspondences from previous loop closures prevent new deformations from breaking the old loop closures. The final cost is optimized using Gauss-Newton gradient descent to retrieve the set of deforming affine transforms \mathcal{G}_T^n . See [3] for more details.

In order to reflect ORB-SLAM loop closures in the dense model in our system using the deformation graph, we first generate a set of surface-to-surface correspondences from

Sequence	$\begin{array}{c} 0 \\ t_{rel}(\%) \end{array}$	$\mathbf{H} \ t_{rel}(\%)$	$\substack{\textbf{H+L}\\t_{rel}(\%)}$	D 3[36] $t_{rel}(\%)$	$\begin{array}{c} \textbf{ORB[20]} \\ t_{rel}(\%) \end{array}$
01	31.17	24.49	25.34	1.07	1.39
02	2.97	3.11	2.38	0.80	0.76
06	1.16	1.42	1.34	0.67	0.51
08	1.90	2.14	2.11	1.00	1.05
09	1.85	2.02	1.62	0.78	0.87
10	1.24	1.50	1.41	0.62	0.60

TABLE I: **Results on test set of KITTI odometry**. Results from ORB-SLAM stereo and a SOTA VO system D3VO are shown for comparison. Though not SOTA, these results show that our system's camera tracking achieves a level of accuracy necessary for building a fused 3D surfel model. It can be seen that using predicted depth leads to accurate RGBD tracking with ORB-SLAM (O). We show the effect of introducing pre-fusion alignment (H) and applying global loop closures to the dense model (H+L).

 $\mathbf{P}_{kf}^{i} \rightarrow \mathbf{\hat{P}}_{kf}^{i}$. To do so, the active surfels $\mathcal{M}^{a} \in \mathcal{M}$ in view of \mathbf{P}_{kf}^{i} are computed, as are the timestamps of the inactive surfels in view of $\mathbf{\hat{P}}_{kf}^{i}$. These timestamps are used during optimization of \mathcal{G}_{T}^{n} to pin the inactive surfels of \mathcal{M} in place. Each of the \mathcal{M}_{s}^{a} generates a constraint consisting of a source point, computed by composing \mathcal{M}_{s}^{a} , with \mathbf{P}_{kf}^{i} and a destination point computed by composing \mathcal{M}_{s}^{a} with $\mathbf{\hat{P}}_{kf}^{i}$. These constraints are then used to optimise the deformation in the same way as [3].

Note that Kf_i is not necessarily the current live camera frame. It could therefore fall outside the current active region of the map. If it is outside the active region of the map, then this will hinder the system's ability to generate surface-tosurface constraints. We found that by setting the active time window appropriately ($\delta t = 200 frames$), the system could find sufficient constraints during loop closures.

Closing loops in this way achieves two goals; adjusting the dense surface geometry to stay consistent with the real world; and keeping the dense map consistent with ORB-SLAM's sparse map and camera pose estimates. It also balances the need for this consistency against the computational intensity of correcting dense geometry. Figure 3 illustrates each of the steps involved in the hybrid loop closure process.

III. EVALUATION

We show quantitative and qualitative results of our system tested on the KITTI odometry benchmark dataset [11]. We have evaluated the accuracy of trajectory estimation and surface reconstruction and provide a breakdown down of the computational performance of the system. All processing was done on a machine with an Intel Core i7 - 7700K CPU, 16GB of RAM, and an NVIDIA GTX 1080 Ti GPU.

A. KITTI - Tracking

The KITTI odometry benchmark dataset provides 11 sequences with ground truth poses. We show results for sequences 01, 02, 06, 08, 09 and 10. The remainder of the sequences are used during training of the depth prediction network and hence ommitted from our evaluation. For each



Fig. 4: Sequence 01 from the KITTI odometry dataset is challenging for monocular systems. (a) and (b) show the live color image from the camera and UnRectDepthNet's depth prediction respectively. (c) and (d) show ElasticFusion's corresponding model predictions. (e) shows a portion of the final global model. (f) shows a comparison between the estimated and ground-truth trajctories.

test sequence we use the relative translational error t_{rel} averaged over sequences ranging in length from 100m to 800m [11]. Table I shows the results alongside a SOTA SLAM system, ORB-SLAM2 [20], and D3VO [36], a SOTA VO system. Interestingly, our experiments show that using the dense depth predictions and the RGBD operating mode of ORB-SLAM (O) results in accurate metric-scale camera tracking with a monocular camera. While introducing hybrid tracking (H) increases error w.r.t the ground truth, it allows the current frame to be fused into the model. Introducing hybrid loop closures (H+L) helps bring the sparse and dense models back into alignment and reduces global error in the model and trajectory. Sequence 01 is challenging for our system. Scenes with little structure lead to a degradation in depth prediction and camera tracking. Figure 4 shows qualitative results of this sequence.

B. KITTI - Surface Reconstruction

We evaluate surface reconstruction accuracy on several sequences from the KITTI odometry benchmark. As KITTI does not include ground-truth surface models, we compare against models constructed from the Velodyne point clouds that accompany each sequence. To do so, we transform each point cloud in the sequence into a global coordinate system using the corresponding ground truth poses. We then apply a voxel filter to downsample the resultant point cloud. In Table II, we show the surface-to-surface mean distance between points in the estimated model and the nearest point in the Velodyne point cloud model. Before computing the score, the two models are rigidly aligned. We use the labels from [44]

Sequence	O (m)	H (m)	H+L (m)
02	5.95	6.40	4.03
06	1.23	0.97	0.64
08	2.70	2.46	2.72
09	1.23	1.18	0.71
10	0.89	0.79	0.82

TABLE II: **Surface Accuracy of our system on KITTI.** Sequence 01 has been ommitted due to poor tracking performance (see Table I)



Fig. 5: **Breakdown of time taken by our system** to process each frame in sequence 09 the KITTI odometry benchmark [11]. Our systems falls just outside of true real-time rates of 10hz for KITTI. In future work we hope to utilize hardware with half precision support to improve system run-time.

Operation	Median(ms)	Mean(ms)	std. dev.(ms)
ORB-SLAM Tracking	40.04	41.93	25.13
Dense Mapping	30.97	35.61	15.82
Depth	60.32	61.14	3.84
Overall	136.48	140.29	31.76

TABLE III: Spread of system run-time over test sequences from the Eigen split of the KITTI odometry dataset

to filter out dynamic objects from Velodyne point clouds.

C. KITTI - Depth Estimation

We use the depth estimation setting on the KITTI Eigen split [32] and report results in Table IV. In our previous work, UnRectDepthNet [30], we outperformed all the previous monocular self-supervised methods. Following best practices, we cap depths at 80 m. We evaluate using the *Improved* [45] ground truth depth maps. We further improve the performance by incorporating additional losses¹ discussed

Method	Abs_{rel}	Sq_{rel}	RMSE	RMSE _{log}
UnRectDepthNet [30]	0.081	0.414	3.412	0.117
+ additional loss terms [46]	0.071	0.334	3.218	0.101
+ improved scale estimation	0.057	0.298	2.95	0.091

TABLE IV: Evaluation of depth estimation on the Improved KITTI Eigen split [45].

in Section II-A. Scale estimation improvements provide an additional improvement in quantitative score, enabling the SLAM pipeline to estimate accurate metric reconstructions.

D. System Resource Usage

In Figure 5 we show a breakdown of the frame processing time for sequence 09 in the KITTI odometry benchmark [11]. Mapping time (blue) goes up in accordance with the number of surfels (purple) as reported in [3]. The system operates between 8 - 9hz. The spike at the end of the sequence is due to a global loop closure. Table III shows the distribution of run-time performance over the test set of the Eigen split of the KITTI odometry dataset.

IV. CONCLUSIONS

We presented a hybrid SLAM system that combines dense depth prediction with sparse feature tracking and dense surfel fusion techniques. The system permits live-dense metric reconstructions of outdoor scenes using a monocular camera in automotive scenarios. Sparse tracking provides camera pose estimation capable of operating robustly at vehicle speeds. The resulting poses are used within the dense fusion tracking step to initialize a whole image alignment refinement process. Global consistency in the model is maintained through visual place recognition and pose to pose constraints from the sparse system, which again are passed to the dense fusion algorithm where they are integrated with a deformation graph-based map correction step. Our results show competitive performance with SOTA techniques while providing dense fused surface models. We believe our system to be the first dense monocular fusion-based visual SLAM system quantitatively evaluated on automotive scenarios. Although the focus of this paper is automotive, the system could easily be adapted to other scenarios by retraining the depth prediction network.

REFERENCES

- [1] R. A. Newcombe, S. Izadi, O. Hilliges, D. Molyneaux, D. Kim, A. J. Davison, P. Kohi, J. Shotton, S. Hodges, and A. Fitzgibbon, "KinectFusion: Real-time dense surface mapping and tracking," in 2011 10th IEEE International Symposium on Mixed and Augmented Reality, 2011, pp. 127–136.
- [2] T. Whelan, M. Kaess, M. Fallon, H. Johannsson, J. Leonard, and J. McDonald, "Kintinuous: Spatially Extended KinectFusion," in RSS Workshop on RGB-D: Advanced Reasoning with Depth Cameras, July 2012.
- [3] T. Whelan, S. Leutenegger, R. S. Moreno, B. Glocker, and A. Davison, "ElasticFusion: Dense SLAM Without A Pose Graph," in *Proc. of Robotics: Science and Systems*, July 2015.
- [4] A. Dai, M. Nießner, M. Zollöfer, S. Izadi, and C. Theobalt, "Bundle-Fusion: Real-time Globally Consistent 3D Reconstruction using Onthe-fly Surface Re-integration," ACM Transactions on Graphics 2017 (TOG), 2017.

¹The loss function improvements will be presented in contemporary work to be published at ICRA 2021.

- [5] R. Mur-Artal, J. M. M. Montiel, and J. D. Tards, "ORB-SLAM: A Versatile and Accurate Monocular SLAM System," *IEEE Transactions* on *Robotics*, vol. 31, no. 5, pp. 1147–1163, 2015.
- [6] C. Campos, R. Elvira, J. Rodríguez, J. Montiel, and J. D. Tard'os, "ORB-SLAM3: An Accurate Open-Source Library for Visual, Visual-Inertial and Multi-Map SLAM," ArXiv, vol. abs/2007.11898, 2020.
- [7] R. Mur-Artal and J. D. Tards, "Visual-Inertial Monocular SLAM With Map Reuse," *IEEE Robotics and Automation Letters*, vol. 2, no. 2, pp. 796–803, 2017.
- [8] J. Engel, T. Schöps, and D. Cremers, "LSD-SLAM: Large-Scale Direct Monocular SLAM," in *Computer Vision – ECCV 2014*, D. Fleet, T. Pajdla, B. Schiele, and T. Tuytelaars, Eds. Springer International Publishing, 2014, pp. 834–849.
- [9] C. Forster, Z. Zhang, M. Gassner, M. Werlberger, and D. Scaramuzza, "SVO: Semidirect visual odometry for monocular and multicamera systems," *IEEE Transactions on Robotics*, vol. 33, no. 2, pp. 249– 265, 2016.
- [10] C. Forster, Z. Zhang, M. Gassner, M. Werlberger, and D. Scaramuzza, "SVO: Semidirect Visual Odometry for Monocular and Multicamera Systems," *IEEE Transactions on Robotics*, vol. 33, no. 2, pp. 249–265, 2017.
- [11] A. Geiger, P. Lenz, and R. Urtasun, "Are we ready for Autonomous Driving? The KITTI Vision Benchmark Suite," in *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2012.
- [12] A. Geiger, P. Lenz, C. Stiller, and R. Urtasun, "Vision meets Robotics: The KITTI Dataset," *International Journal of Robotics Research* (*IJRR*), 2013.
- [13] Y. Wang, W.-L. Chao, D. Garg, B. Hariharan, M. Campbell, and K. Q. Weinberger, "Pseudo-LiDAR From Visual Depth Estimation: Bridging the Gap in 3D Object Detection for Autonomous Driving," in 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 2019, pp. 8437–8445.
- [14] H. Strasdat, J. M. M. Montiel, and A. J. Davison, "Scale drift-aware large scale monocular slam," in *Proc. Robotics: Science and Systems* (*RSS*), 2010.
- [15] R. A. Newcombe, S. J. Lovegrove, and A. J. Davison, "DTAM: Dense tracking and mapping in real-time," in 2011 International Conference on Computer Vision, 2011, pp. 2320–2327.
- [16] J. Engel, J. Sturm, and D. Cremers, "Semi-dense Visual Odometry for a Monocular Camera," in *Proc. of IEEE International Conference on Computer Vision (ICCV)*, Dec. 2013.
- [17] A. J. Davison, I. D. Reid, N. D. Molton, and O. Stasse, "MonoSLAM: Real-Time Single Camera SLAM," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 29, no. 6, pp. 1052–1067, 2007.
- [18] G. Klein and D. Murray, "Parallel Tracking and Mapping for Small AR Workspaces," in *Proc. Sixth IEEE and ACM International Symposium* on *Mixed and Augmented Reality*, Nov. 2007.
- [19] E. Eade and T. Drummond, "Unified loop closing and recovery for real time monocular slam," in *In British Machine Vision Conference*, 2008.
- [20] R. Mur-Artal and J. D. Tards, "ORB-SLAM2: An Open-Source SLAM System for Monocular, Stereo, and RGB-D Cameras," *IEEE Transactions on Robotics*, vol. 33, no. 5, pp. 1255–1262, 2017.
- [21] T. Dang, C. Hoffmann, and C. Stiller, "Continuous Stereo Self-Calibration by Camera Parameter Tracking," *IEEE Transactions on Image Processing*, vol. 18, no. 7, pp. 1536–1550, 2009.
- [22] J. Abel and J. Chaffee, "Existence and uniqueness of GPS solutions," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 27, no. 6, pp. 952–956, 1991.
- [23] J. Engel, V. Koltun, and D. Cremers, "Direct sparse odometry," *IEEE transactions on pattern analysis and machine intelligence*, vol. 40, no. 3, pp. 611–625, 2017.
- [24] X. Gao, R. Wang, N. Demmel, and D. Cremers, "LDSO: Direct Sparse Odometry with Loop Closure," in *iros*, Oct. 2018.
- [25] T. Schps, T. Sattler, and M. Pollefeys, "BAD SLAM: Bundle Adjusted Direct RGB-D SLAM," in 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 2019, pp. 134–144.
- [26] M. F. Fallon, P. Marion, R. Deits, T. Whelan, M. Antone, J. McDonald, and R. Tedrake, "Continuous humanoid locomotion over uneven terrain using stereo fusion," in 2015 IEEE-RAS 15th International Conference on Humanoid Robots (Humanoids), 2015, pp. 881–888.
- [27] M. Bloesch, J. Czarnowski, R. Clark, S. Leutenegger, and A. J. Davison, "CodeSLAM Learning a Compact, Optimisable Representation

for Dense Visual SLAM," in *Proc. of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.

- [28] E. Sucar, S. Liu, J. Ortiz, and A. Davison, "iMAP: Implicit Mapping and Positioning in Real-Time," in arXiv preprint arXiv:2103.12352, 2021.
- [29] G. Costante, M. Mancini, P. Valigi, and T. A. Ciarfuglia, "Exploring Representation Learning With CNNs for Frame-to-Frame Ego-Motion Estimation," *IEEE Robotics and Automation Letters*, vol. 1, no. 1, pp. 18–25, 2016.
- [30] V. Ravi Kumar, S. Yogamani, M. Bach, C. Witt, S. Milz, and P. Mäder, "UnRectDepthNet: Self-Supervised Monocular Depth Estimation using a Generic Framework for Handling Common Camera Distortion Models," in *IEEE/RSJ International Conference on Intelligent Robots* and Systems, IROS 2020, 2021, 2020, pp. 8177–8183.
- [31] T. Zhou, M. Brown, N. Snavely, and D. G. Lowe, "Unsupervised Learning of Depth and Ego-Motion from Video," in *IEEE Conference* on Computer Vision and Pattern Recognition (CVPR), 2017, pp. 6612– 6619.
- [32] D. Eigen, C. Puhrsch, and R. Fergus, "Depth Map Prediction from a Single Image Using a Multi-Scale Deep Network," in *Proc. of 27th International Conference on Neural Information Processing Systems -Volume 2*, ser. NIPS'14, 2014, p. 23662374.
- [33] C. Godard, O. Mac Aodha, and G. J. Brostow, "Unsupervised Monocular Depth Estimation With Left-Right Consistency," in *Proc. of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017.
- [34] J. Czarnowski, T. Laidlow, R. Clark, and A. Davison, "Deepfactors: Real-time probabilistic dense monocular slam," *IEEE Robotics and Automation Letters*, vol. 5, pp. 721–728, 2020.
- [35] K. Tateno, F. Tombari, I. Laina, and N. Navab, "CNN-SLAM: Real-Time Dense Monocular SLAM with Learned Depth Prediction," in 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2017, pp. 6565–6574.
- [36] N. Yang, L. von Stumberg, R. Wang, and D. Cremers, "D3VO: Deep Depth, Deep Pose and Deep Uncertainty for Monocular Visual Odometry," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [37] H. Zhou, B. Ummenhofer, and T. Brox, "DeepTAM: Deep Tracking and Mapping," in *Proc of European Conference on Computer Vision* (ECCV), Sept. 2018.
- [38] S. Wang, R. Clark, H. Wen, and N. Trigoni, "Deepvo: Towards end-to-end visual odometry with deep recurrent convolutional neural networks," in 2017 IEEE International Conference on Robotics and Automation (ICRA), 2017, pp. 2043–2050.
- [39] K. Wang, F. Gao, and S. Shen, "Real-time scalable dense surfel mapping," in *International Conference on Robotics and Automation* (*ICRA*), 2019, pp. 6919–6925.
- [40] C. Shu, K. Yu, Z. Duan, and K. Yang, "Feature-metric loss for selfsupervised learning of depth and egomotion," in *Proc. of European Conference on Computer Vision (ECCV)*. Springer, 2020, pp. 572– 588.
- [41] R. Garg, V. K. Bg, G. Carneiro, and I. Reid, "Unsupervised cnn for single view depth estimation: Geometry to the rescue," in *Proc. of European Conference on Computer Vision (ECCV)*. Springer, 2016, pp. 740–756.
- [42] T. Zhou, M. Brown, N. Snavely, and D. G. Lowe, "Unsupervised learning of depth and ego-motion from video," in *Proc. of IEEE* conference on Computer Vision and Pattern Recognition (CVPR), 2017, pp. 1851–1858.
- [43] D. Galvez-Lpez and J. D. Tardos, "Bags of binary words for fast place recognition in image sequences," *IEEE Transactions on Robotics*, vol. 28, no. 5, pp. 1188–1197, 2012.
- [44] J. Behley, M. Garbade, A. Milioto, J. Quenzel, S. Behnke, C. Stachniss, and J. Gall, "SemanticKITTI: A Dataset for Semantic Scene Understanding of LiDAR Sequences," in *Proc. of the IEEE/CVF International Conf. on Computer Vision (ICCV)*, 2019.
- [45] J. Uhrig, N. Schneider, L. Schneider, U. Franke, T. Brox, and A. Geiger, "Sparsity Invariant CNNs," in *Proc. of 3DV*, Oct. 2017, pp. 11–20.
- [46] V. R. Kumar, S. Yogamani, H. Rashed, G. Sitsu, C. Witt, I. Leang, S. Milz, and P. Mäder, "Omnidet: Surround view cameras based multi-task visual perception network for autonomous driving," *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 2830–2837, 2021.