

Neural Semantic Map-Learning for Autonomous Vehicles

Markus Herb^{1,2}, Nassir Navab¹ and Federico Tombari^{1,3}

Abstract—Autonomous vehicles demand detailed maps to maneuver reliably through traffic, which need to be kept up-to-date to ensure a safe operation. A promising way to adapt the maps to the ever-changing road-network is to use crowd-sourced data from a fleet of vehicles. In this work, we present a mapping system that fuses local submaps gathered from a fleet of vehicles at a central instance to produce a coherent map of the road environment including drivable area, lane markings, poles, obstacles and more as a 3D mesh. Each vehicle contributes locally reconstructed submaps as lightweight meshes, making our method applicable to a wide range of reconstruction methods and sensor modalities. Our method jointly aligns and merges the noisy and incomplete local submaps using a scene-specific Neural Signed Distance Field, which is supervised using the submap meshes to predict a fused environment representation. We leverage memory-efficient sparse feature-grids to scale to large areas and introduce a confidence score to model uncertainty in scene reconstruction. Our approach is evaluated on two datasets with different local mapping methods, showing improved pose alignment and reconstruction over existing methods. Additionally, we demonstrate the benefit of multi-session mapping and examine the required amount of data to enable high-fidelity map learning for autonomous vehicles.

I. INTRODUCTION

Maps are crucial for many applications such as mobile robots including autonomous vehicles or augmented reality to facilitate localization, navigation and interaction with the environment. While maps can prove beneficial for higher-level tasks, it is paramount that the map content is accurate and up-to-date, especially in safety critical settings such as autonomous driving. Traditionally, dedicated mapping vehicles have been used for map data collection, which however quickly becomes infeasible to cover areas beyond city-scale with a high update rate due to the limited number of vehicles and large operating effort. Collecting crowd-sourced user data from car fleets is a promising alternative to gather fresh mapping data in a timely manner [1]–[4]. In contrast to specialized mapping vehicles, commodity cars typically come with a relatively limited sensor set, storage and upload capacities, which demands efficient use of resources to be worthwhile for mapping large scale road networks.

To address these challenges, we present a centralized mapping system producing dense semantic maps, depicted in Fig. 1, by incorporating map contributions from many individual vehicles or *agents*. We chose to represent the individual submaps as 3D semantic surface meshes, as these can represent dense geometry and semantics without the




Fig. 1: Our neural multi-session mapping approach fuses dense submaps created from different agents through local dense SLAM by training a neural field to predict dense semantics and geometry that best explains the submaps collected from multiple mapping sessions.

need for spatial discretization, are lightweight in terms of storage demands and can be produced using various mapping approaches relying on different sensor configurations. Furthermore, the meshes can often be easily extracted from the local environment representation each agent builds for other tasks such as obstacle avoidance, thus saving computational effort by fusing the local sensor data only once. Our proposed system hence avoids the need for raw sensor data to be sent to a central server altogether and instead uses already fused submaps as intermediate mapping data uploaded by the agents to the server. While being friendly for low-bandwidth mobile connections, this also reduces privacy concerns associated with image data.

Our method builds on recent advances in Neural Implicit Representations and uses a Neural Signed Distance Field [5] to represent geometry and semantics of the environment. During reconstruction, we supervise the neural field to produce a signed distance and semantic logits for each point in space that best matches all the input submaps observed by different agents exploring the scene. Since the local maps produced by each vehicle are only coarsely aligned using

¹Technical University of Munich, Computer-Aided Medical Procedures, Garching bei München, Germany

²CARIAD SE, Ingolstadt, Germany

³Google, Zurich, Switzerland

inaccurate GPS information, we also optimize the relative poses of all submaps jointly during the reconstruction. Our fundamental assumption is that neural fields are capable to learn a *fused, coherent* scene representation from noisy, incomplete and potentially contradicting input submaps observed from different agents. To enable fast update times and scalability to large regions, we cluster the mapped areas into independent tiles, each represented by a feature OcTree [6] or HashGrid [7] as efficient spatial grid-structure with prediction heads for semantics and geometry.

To summarize, our main contributions in this work are

- a novel multi-session mapping system for dense semantic submap fusion using neural fields,
- extensive experiments on two real-world datasets demonstrating the efficacy of our method for crowd-sourced map learning

II. RELATED WORK

Mapping has been a central topic in the robotics community for decades, with an abundance of work in the realm of Simultaneous Localization and Mapping (SLAM). We focus our review on the most related works in the fields of maps for autonomous vehicles, dense reconstruction and multi-robot mapping using vision sensors.

A. Mapping for Autonomous Vehicles

Crowd sourced mapping has seen a surge in interest in recent years, particularly for the application of autonomous vehicles. However, most works in the field focus on specific map features related to road networks, such as lane markings or traffic signs [1], which are light-weight and easy to gather in a crowd-sourced fashion. A special focus in these works has been put on detailed mapping of the geometry [3], [8], topology [2] or dealing with change detection and update [4], [9] for lane markings. However, these works focus only on specific road environments and offer a limited spatial understanding due to the extreme level of environment abstraction, which restricts its usefulness for tasks like localization or augmented reality. In contrast, Cheng *et al.* [10] use semantically labeled sparse points from visual odometry that are more generic, but still sparse.

B. Dense Map Reconstruction

Structure-from-Motion (SfM) methods such as COLMAP [11] are well established for large-scale sparse and dense reconstruction from image-collections. While SfM methods produce high-quality dense reconstructions, they demand significant computational resources for multi-view stereo and need to collect large amounts of image data. As an alternative, fast online reconstructions using surfels [12] or meshes [13] have been proposed for robotic applications.

While the aforementioned approaches represent 3D structure *explicitly* as points, meshes or surfels, implicit methods leveraging Signed Distance Fields (SDF) as representation have been used for high-quality 3D reconstructions. Hashed voxel-grid structures are commonly used to reconstruct an

SDF [14], especially for small-scale environments, with extensions towards larger-scale [15] and semantic information [16] enabling map reconstructions for street-scale maps [17]. However, very detailed large-scale outdoor reconstructions using voxel-representations are limited by prohibitively high memory requirements arising from small voxel-sizes.

Recently, Neural Implicit Representations pioneered by NeRFs [18] for view synthesis and DeepSDF [5] for geometry reconstruction have shown groundbreaking results by learning properties of a specific scene such as density, color or signed distance as a Multi-Layer Perceptron (MLP) that can be queried for every 3D point in the scene. Early works on large-scale scene reconstruction using images [19] and LiDAR [20] show impressive results, but are slow to train. Using hashed feature-grids [7] or OcTrees [6], [21] for storing scene features in combination with a shallow MLP head has enabled dramatic speedups for neural field training and inference, especially on large-scale scenes. In addition to just reconstructing the scene content, joint optimization with input poses for cameras [22] and LiDARs [23] has been proposed to overcome the need for posed sensor data.

C. Multi-Robot Mapping

While a majority of SLAM works address the problem of single-robot mapping, many practical applications require the use of multiple agents working together in a centralized or decentralized setup to map larger environments or keep them up to date.

ORB_SLAM-Atlas [24] and CCM-SLAM [25] both build on top of ORB-SLAM [26] to extend it to a centralized multi-robot mapping system, with a focus on map accuracy and inter-agent communication respectively. `maplab 2.0` [27] is another framework for multi-robot multi-session mapping that follows a centralized approach, which supports various modalities including visual keypoints, inertial and GNSS data. However, due to its sparse map representation, all of these approaches are primarily suited to build maps for (re-)localization tasks.

Multiple works have expanded the multi-robot SLAM problem to incorporate semantic and dense map information in various ways. Kimera-Multi [28] extends Kimera [16] into a decentralized mapping system, however the multi-robot part focuses on a joint sparse pose-graph optimization, while each agent keeps and refines its dense mesh reconstruction, such that no dense map information is created collaboratively.

Golodetz *et al.* [29] use a centralized approach, where agents transmit RGB-D data to a powerful server that performs map reconstruction for all robots. In [30] a decentralized multi-robot mapping system is presented that uses stereo cameras to reconstruct TSDF-based submaps that are shared among robots and geometrically aligned as part of a pose-graph optimization. Both such approaches are hard to scale to large fleets of robots due to large amount of resources requirement for data exchange and storage. Coxgraph [31] uses a similar approach, but compresses the TSDF maps into meshes for transmission to a central mapping server,

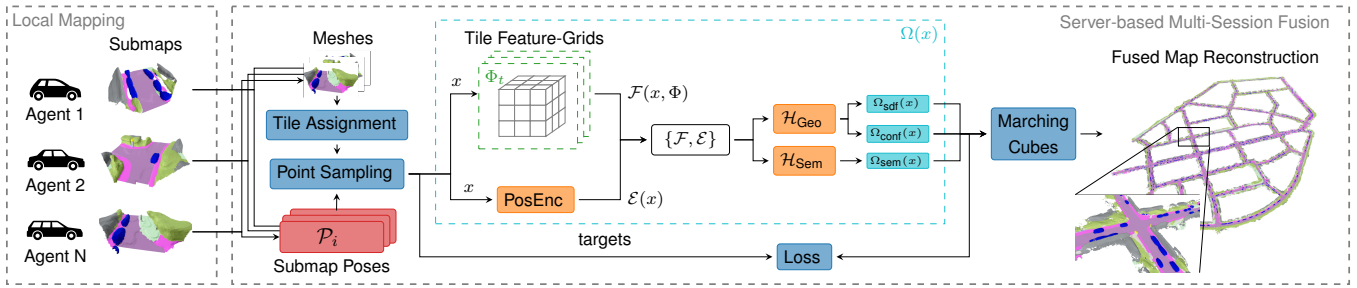


Fig. 2: System Pipeline overview of our method: We collect meshes and poses \mathcal{P} for each local agent mapping session, cluster them into different geographic tiles t , and use them to supervise the neural semantic signed distance field. This consists of the feature grids Φ_t for each tile and shared geometry head \mathcal{H}_{Geo} predicting SDF $s(x) = \Omega_{\text{sdf}}(x)$ and confidence $c(x) = \Omega_{\text{conf}}(x)$ and semantic head \mathcal{H}_{Sem} predicting logits $l(x) = \Omega_{\text{sem}}(x)$, conditioned on grid features \mathcal{F} and point embeddings \mathcal{E} . We jointly optimize submap poses \mathcal{P}_i , neural field grids Φ_t and decoder heads \mathcal{H} . The final fused map reconstruction can be conveniently extracted from the neural field using Marching Cubes.

where the SDF is recovered from the mesh. We also adopt meshes as exchange format for map data from local agents to a centralized server in our system, enabling lightweight data transmission and compatibility with a wide variety of local map reconstruction methods. In contrast to existing approaches, we employ a neural implicit representation to fuse the map data contributed from each agent. This allows us to optimize geometry, semantics and poses jointly without the need for specific localization landmarks and facilitate large-scale high-fidelity map reconstruction.

III. APPROACH

A. System Overview

In Fig. 2 we show an overview of our pipeline, which is made up of the local map generation in each agent, which we describe in section III-B, and the main focus of our work, the fusion component, explained in section III-C. In our system, each mapping agent produces local submaps from onboard sensory data, which are uploaded to a central mapping server by each agent as a semantically labelled mesh. Afterwards, we use all submaps to supervise the neural field to produce a fused semantic signed distance field.

B. Submap Generation

On each agent, we run a local dense semantic SLAM to generate mapping data for the environment that the agent explores. In this work we focus on crowd-sourced mapping using low-cost monocular-only cameras, but our method is applicable to any dense semantic SLAM system, e.g. also LiDAR based methods. We use ORB-SLAM2 [32] in monocular configuration to generate local vehicle odometry for each agent. To demonstrate the applicability to different local mapping systems, we show results using a lightweight mesh mapping from semantic image segmentation and monocular depth [13] as well as single RGB semantic scene completion [33] for dense mesh reconstruction. In sum, each agent transmits local vehicle odometry as well as corresponding GPS measurements and a number of submaps encoded as triangle mesh, with vertices, triangles, and semantic labels for each face or vertex.

C. Neural Map Fusion

After collecting all submaps from each agent-session, we first cluster all submaps into fixed-size square tiles using the provided coarse GPS measurements. This ensures scalability to larger scenes as only a subset of tiles-of-interest can be loaded for training or inference and is comparable to voxel-hashing techniques in traditional voxel-based surface reconstruction.

1) *Neural Semantic Signed Distance Field*: We represent the fused map as a neural field $\Omega(x)$, a neural network which learns to predict the geometry and semantics for arbitrary input sampling points $x \in \mathbb{R}^3$. As in other neural surface reconstruction methods [21], [23], [34], we use a signed distance field denoted as $s(x) = \Omega_{\text{sdf}}(x) : \mathbb{R}^3 \mapsto \mathbb{R}$ to model the geometry, where the zero-levelset defines the surface implicitly, from which a dense surface mesh can be extracted by applying marching cubes on a regularly sampled grid. We use dedicated feature-grids Φ_t for each spatial tile t and combine these with shallow multi-layer perceptrons (MLPs) as output heads \mathcal{H} shared across all tiles, which ensures scalability while keeping training time low. In a first step, we interpolate point-wise features $\mathcal{F}(x, \Phi_t)$ tri-linearly from the corresponding grid, which are stacked with the positional encodings $\mathcal{E}(x)$ of the input coordinates x and feed them into the prediction heads $\mathcal{H}_{\text{head}}(\{\mathcal{F}(x, \Phi_t), \mathcal{E}(x)\}) = \{\Omega_{\text{pred}}(x), \dots\}$. While positional encoding is not commonly used in existing neural fields conditioned on 3D grid features [7], [21], we found that it improves the pose optimization (see section IV-G.1).

a) *Feature Grids*: We adopt OcTrees [6] and HashGrids [7] as sparse and memory-efficient feature-grid structures, which are commonly used in neural fields to enable large-scale reconstruction. For both grids, features are stored in a lookup-table (codebook) with the mapping of point coordinates to features being defined by either the voxel-index inside the OcTree or by a spatial hash function. We evaluate the suitability of both structures for our fusion approach in section IV.

b) *Geometry Head*: Our geometry MLP head $\mathcal{H}_{\text{Geo}}(x) = \{s(x), c(x)\}$ primarily predicts the signed

distance field $s(x) = \Omega_{\text{sdf}}(x)$. In addition, it outputs a confidence score $c(x) = \Omega_{\text{conf}}(x) : \mathbb{R}^3 \mapsto [0, 1]$, which represents the likelihood of the surface or objects existence. While the signed distance in an SDF represents the estimated distance towards the closest object surface, there is no notion of existence measure associated to it. When using data from multiple, contradictory sessions with noisy or spurious measurements, it is crucial to denote how reliable a surface prediction is. For instance, some temporary object might be detected only by one or few agents and this should be modelled accordingly by the neural field.

c) Semantic Head: In addition to the geometry head, we use a dedicated semantic head $\mathcal{H}_{\text{Geo}}(x) = \{l(x)\}$, which predicts semantic logits $l(x) = \Omega_{\text{sem}}(x) : \mathbb{R}^3 \mapsto \mathbb{R}^K$. From this, we derive class scores $L(x_{\text{surf}}) = \sigma(l(x_{\text{surf}})) \in [0, 1]^K$ with σ being the *softmax* operation. We found that using a dedicated decoder MLP for semantics slightly improved the reconstruction quality.

2) Pose Optimization: We expect the input submaps contributed by each agent to be geo-localized, as GPS sensors are commonly available. However, the typical error of GPS measurements is relatively large compared to the required accuracy needed for centimeter-accurate mapping. In order to overcome this, we optimize both the neural field parameters and the input poses jointly to produce optimal reconstruction results. For each input submap i , we estimate a pose \mathcal{P}_i in global coordinates, such that local submap coordinates x_i can be transformed to global coordinates $\tilde{x}_i = \mathcal{T}(\mathcal{P}_i, x_i)$. Each pose $\mathcal{P}_i \in \mathfrak{se}(3)$ is represented as Lie algebra to facilitate optimization and is initialized with the GPS measurement $\hat{\mathcal{P}}_{\text{GPS}, i}$.

D. Map Reconstruction & Alignment

1) Supervision Point Sampling: To fuse the submaps into a common environment reconstruction, we train the neural field using individual points, sampled from the input meshes. For each batch, we sample points x_i uniformly distributed on the surface of the input meshes i and transform them to the global frame using corresponding poses $\tilde{x}_i = \mathcal{T}(\mathcal{P}_i, x_i)$ to obtain surface samples $\hat{x}_{\text{surf}} = \{\tilde{x}_0, \tilde{x}_1, \dots\}$. From these we compute target samples $x_{\text{surf}} = \hat{x}_{\text{surf}} + \hat{s} \cdot \hat{n}_{\text{surf}}$ around the surface at an offset $\hat{s} \sim \mathcal{N}(0, \sigma_s)$ with $\sigma_s = 0.05m$ along the surface normal \hat{n}_{surf} such that \hat{s} represents the target signed distance for x_{surf} . We also sample the same number of points x_{space} uniform randomly in all input submap bounding boxes for regularization and confidence priors.

2) Losses: Our primary objective is the surface reconstruction, for which we regress SDF value of the target points x_{surf} sampled around the surface

$$\mathcal{L}_{\text{sdf}}(s(x_{\text{surf}})) = (s(x_{\text{surf}}) - \hat{s})^2 \quad (1)$$

with \hat{s} being the offset distance of the sampled point along the surface normal. We further regress the predicted surface normals $n(x_{\text{surf}})$ against the normals sampled from the input meshes \hat{n}_{surf} for each point by computing the normal prediction as $n(x) = \frac{\partial}{\partial x} s(x)$ using the numerical gradient method

from [34] and formulate the loss as

$$\mathcal{L}_{\text{norm}}(n(x_{\text{surf}})) = \|n(x_{\text{surf}}) - \hat{n}_{\text{surf}}\|^2. \quad (2)$$

For semantics, we use a standard cross-entropy (CE) loss on predicted class scores $L(x_{\text{surf}}) \in [0, 1]^K$ and target labels $\hat{L}_{\text{surf}} \in \mathbb{N}$:

$$\mathcal{L}_{\text{sem}}(L(x_{\text{surf}})) = \text{CE}(L(x_{\text{surf}}), \hat{L}_{\text{surf}}). \quad (3)$$

The target labels \hat{L}_{surf} for each point x_{surf} sampled around the surface correspond to the label of the input mesh face the surface point was sampled from. To regularize the SDF, we add an eikonal loss [35] to all sampled surface and volume points:

$$\mathcal{L}_{\text{eik}}(x_{\text{surf}, \text{space}}) = \left(\left\| \frac{\partial}{\partial x_{\text{surf}, \text{space}}} s(x_{\text{surf}, \text{space}}) \right\|_2 - 1 \right)^2 \quad (4)$$

enforcing the gradient-vector of the SDF w.r.t. input coordinates to be close to unit-norm. We supervise the confidence output using a binary cross-entropy (BCE) loss to predict 1 for valid surfaces x_{surf} and 0 for non-surface areas x_{space} :

$$\mathcal{L}_{\text{conf}}(c(x_{\text{surf}}), c(x_{\text{space}})) = \text{BCE}(c(x_{\text{surf}}), 1) + \text{BCE}(c(x_{\text{space}}), 0) \quad (5)$$

We add an odometry loss to stabilize the pose optimization by enforcing that the relative pose $\mathcal{P}_{i \rightarrow i+1} = \mathcal{P}_{i+1}^{-1} \mathcal{P}_i$ of consecutive submaps $i, i+1$ is consistent with their measured odometry $\hat{\mathcal{P}}_{\text{odom}, i \rightarrow i+1}$

$$\mathcal{L}_{\text{odom}}(\mathcal{P}_{i \rightarrow i+1}) = w_{\text{odom}} \hat{\mathcal{P}}_{\text{odom}, i \rightarrow i+1}^{-1} \mathcal{P}_{i \rightarrow i+1} \quad (6)$$

where w_{odom} is weight-vector to model the uncertainty in the odometry. Our total loss is therefore composed as

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{sdf}} + \lambda_{\text{norm}} \mathcal{L}_{\text{norm}} + \lambda_{\text{sem}} \mathcal{L}_{\text{sem}} + \lambda_{\text{eik}} \mathcal{L}_{\text{eik}} + \lambda_{\text{conf}} \mathcal{L}_{\text{conf}} + \lambda_{\text{odom}} \mathcal{L}_{\text{odom}} \quad (7)$$

IV. EVALUATION

In the following, we present an evaluation of our approach on two different datasets using different methods for submap-mesh generation.

A. Experimental Setup

We implemented the proposed approach in PyTorch using the feature HashGrid and OcTree implementation of Kaolin Wisp [36], with modifications to ensure gradient propagation for pose optimization in PyTorch. We train all tiles of size $128m \times 128m$ jointly, randomly selecting one tile per iteration, using AdamW with learning rate 10^{-2} , weight decay of 10^{-2} and LR decay of $10^{-0.001 \frac{\text{iter}}{\#\text{tiles}}}$ for grid and model parameters. Pose parameters use learning rates of 10^{-2} and 10^{-4} for translational and rotational components respectively, without weight decay or LR decay. We train for an average of 500 iterations per tile with 125000 surface and volume samples each per batch. For the HashGrid, we use 16 levels of feature size 2, with grid sizes from 2^4 to 2^{11} and a codebook size of $2^{N_c} = 2^{16}$. Our OcTree

configuration uses 4 feature-levels with leaf-voxel size of $0.5m$ and feature size $N_F = 4$ per level. We further apply progressive level-of-detail (LOD) optimization [34], where we activate only the coarser half of feature-levels in the beginning and incrementally add one additional layer every L iterations. Inactive levels are masked out to zero. We use separate MLP heads for geometry and semantics with 2 hidden layers and 128 neurons each. All experiments were conducted on an NVIDIA V100 with 32GB VRAM, with training taking approximately 90 seconds per tile. We use $\lambda_{\text{norm}} = \lambda_{\text{sem}} = \lambda_{\text{conf}} = \lambda_{\text{odom}} = 1, \lambda_{\text{eik}} = 0.1$ and a confidence threshold c_{th} of 0.7 when extracting meshes for all datasets.

B. Datasets

a) *KITTI360*: KITTI360 [37] is an outdoor driving dataset covering suburban areas with front-facing stereo camera and semantically labeled LiDAR point clouds as ground truth. While the dataset only contains single traversals of the same scene, the trajectories are very loopy, often covering the same area in forward and backward direction. To further increase the coverage, we use only monocular camera data and treat left and right stereo camera as independent traversals. To generate semantic input meshes to be fused, we use MonoScene [33] to predict a semantic occupancy map, which we convert into a triangle mesh. We choose sequences 6 and 9 for evaluation and use the rest to fine-train MonoScene on KITTI360, initialized from the official model trained on SemanticKITTI. We use the SLAM-optimized frame poses as ground-truth and the raw OXTS measurements as GPS poses, which we disturb using artificial gaussian noise since OXTS already provides highly accurate GPS poses. We use every 5th keyframe for mapping and group consecutive frames into chunks of length 5.

b) *HD Map Dataset*: In addition, we use a private dataset that features 7 traversals of an urban route of around 3 km in length in forward and backward direction using a single front-facing camera mounted in the vehicle and standard GPS. As ground truth, we use a manually annotated high-definition map and global pose measurements from RTK-GPS. We use the method in [13] to generate lightweight mesh submaps of 10 keyframes each.

C. Baseline

We compare our system against traditional map-alignment methods ICP [38] and Generalized ICP [39], which we apply in a factor-graph with additional pose constraints as in our approach. For each submap, we select 3 neighboring submaps with the highest spatial overlap. On each submap-mesh we sample 500 points uniformly and compute the (G)ICP-loss over all submap-pairings. We additionally implement a semantic-mode that only associates points with matching semantic label. We optimize the alignment loss with robust Huber kernel using the same optimizer as our neural field, and use the aligned poses as input for the map reconstruction.

For map reconstruction, we build a simple but strong baseline by applying state-of-the-art mesh reconstruction from oriented points using Poisson Surface Reconstruction [40] in COLMAP [11] on 1000000 points with surface normals sampled from all input submaps for each tile. We found that using more sample points did not improve the reconstruction while increasing processing time significantly. After mesh-fusion, we transfer labels from the input meshes by selecting the most-common label of the 5 closest points to each mesh face.

D. Evaluation Metrics

Map Reconstruction: To evaluate the quality of the reconstructed map, we adopt precision and recall metrics from [41], measuring accuracy and completeness of the reconstruction respectively, and the F_1 -Score as harmonic mean of both as geometric (*geo*) reconstruction score. We additionally compute the corresponding semantic (*sem*) scores as average over per-class F-scores. We choose a distance threshold of $d = 20cm$.

Pose Alignment: For pose optimization, we report the absolute root-mean squared (RMSE) translation and rotation errors between estimated poses and corresponding ground-truth poses after rigid alignment of the entire trajectory to the ground-truth.

E. Quantitative Results

In Table I we provide an overview of alignment and reconstructions results for different combinations of pose alignment and geometry reconstruction. We notice that our proposed method, which optimizes poses and reconstruction jointly, achieves significantly better alignment and reconstruction results over the baseline across all datasets. For the ICP-based alignment, we found that point-to-point ICP and generalized ICP perform overall similarly with slight benefits for either on the different datasets. Including semantic information has only marginal effect on the alignment. In our neural field reconstruction, we find that HashGrid and OcTree representations achieve overall similar alignment and reconstruction results.

F. Qualitative Results

Fig. 3 shows qualitative reconstruction results of our approach. We compare neural reconstructions using noisy GPS, optimized alignment and oracle (GT) poses as well as Poisson reconstruction. While direct reconstruction from noisy GPS leads to inferior results, joint neural alignment and reconstruction and factor-graph aligned Poisson meshing achieve visual quality close to reconstruction with oracle poses. We note that Poisson reconstruction lacks some detail in the reconstruction, e.g. road markings, while neural reconstruction is crisper. We present additional qualitative results in our supplementary video.

G. Ablation Studies

In the following we present multiple ablation studies to demonstrate the effect of individual components of our system.

TABLE I: Alignment and reconstruction results overview for different combinations of pose-alignment and map reconstructions. Scores represent translational [m] and rotational [deg] RMSE as well as geometric (geo) and mean semantic (sem) F-Score for all datasets. Highlighted results denote **best** and *second-best* per column.

Method	HD-Map				Kitti360 06				Kitti360 09				
	trans ↓	rot ↓	geo ↑	sem ↑	trans ↓	rot ↓	geo ↑	sem ↑	trans ↓	rot ↓	geo ↑	sem ↑	
Poisson	ICP	1.029	1.104	0.333	0.152	3.774	3.466	0.108	0.051	4.044	3.485	0.089	0.047
	ICP+Sem	1.016	1.101	0.361	0.169	3.818	3.513	0.107	0.049	3.997	3.524	0.086	0.053
	GICP	1.256	1.432	0.314	0.148	2.197	3.685	0.100	0.059	2.039	3.608	0.127	0.084
	GICP+Sem	1.227	1.425	0.277	0.138	2.163	3.631	0.103	0.060	2.086	3.618	0.125	0.091
Neural Field (ours)	OcTree	<i>0.830</i>	<i>1.030</i>	<i>0.552</i>	<i>0.273</i>	<i>1.710</i>	2.040	0.168	0.104	1.057	1.568	0.431	0.280
	HashGrid	0.810	0.924	0.605	0.294	1.603	<i>2.069</i>	<i>0.159</i>	<i>0.099</i>	<i>1.076</i>	<i>1.597</i>	<i>0.410</i>	<i>0.266</i>

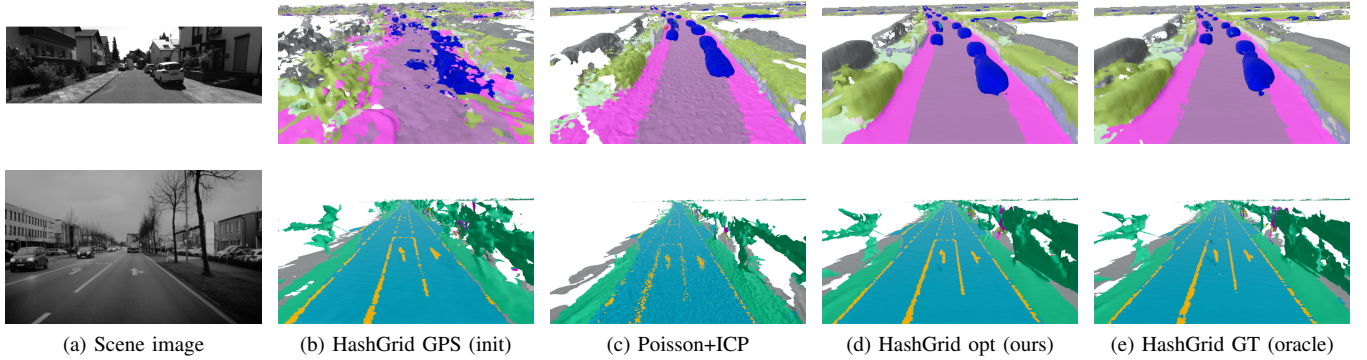


Fig. 3: Qualitative reconstruction results for KITTI360 (top) and HD-Map (bottom).

1) *Submap Alignment*: In Table II we demonstrate the effect of different alignment configurations. We observe that incorporating relative pose constraints is important to stabilize convergence. Furthermore, both the added positional encoding and the incremental level-of-detail reduce the alignment error individually for HashGrids and result in the best overall configuration when used jointly for both feature representations.

TABLE II: Alignment results on HD-Map dataset for neural field (NeF) optimization with relative pose constraints (rel), positional encoding (pe) and progressive level-of-detail (lod). Scores represent mean±std. of translational [m] and rotational [deg] RMSE over 3 runs.

Pose Opt.				HashGrid		OcTree	
align	rel	pe	lod	trans ↓	rot ↓	trans ↓	rot ↓
none				1.42	2.15	1.42	2.15
				1.21±0.01	2.03±0.01	1.13±0.01	1.93±0.02
	✓			1.06±0.02	1.32±0.03	0.96±0.00	0.99±0.02
NeF	✓	✓		0.97±0.01	1.08±0.01	0.88±0.00	1.04±0.00
	✓		✓	0.85±0.01	1.14±0.01	0.88±0.01	0.94±0.01
	✓	✓	✓	0.81±0.00	0.94±0.01	0.84±0.01	1.02±0.02

2) *Reconstruction*: In addition to joint alignment and reconstruction, we also evaluate the reconstruction results with various fixed input poses using GPS (input), aligned using neural optimization and ground-truth poses in Table III. We observe that all reconstruction methods perform similarly for semantic scores, while Poisson reconstruction achieves

slightly higher scores in the geometric reconstruction, as long as input poses are of high quality.

TABLE III: Reconstruction scores for different reconstruction methods for fixed submap poses from varying sources on HD-Map dataset. *Aligned* denotes reconstruction from fixed input poses pre-optimized using Neural Field optimization.

Recon	GPS		Aligned		Ground-Truth	
	geo ↑	sem ↑	geo ↑	sem ↑	geo ↑	sem ↑
Poisson	0.140	0.067	0.636	0.329	0.737	0.408
OcTree	0.177	0.086	0.610	0.309	0.712	0.376
HashGrid	0.122	0.063	0.613	0.300	0.721	0.392

3) *Multi-Session Mapping*: In this section, we provide experimental results on our HD-Map dataset with varying number of total mapping sessions as well as with different submap sizes, which impacts the amount of data that needs to be collected from mapping agents. We ran our reconstruction for each number of sessions $N \in [1, 7]$ with 3 randomly sampled sets each from the entire dataset for both forward and backward direction. In Fig. 4a we can observe that with an increasing number of mapping sessions the reconstruction quality rises as well, supporting our motivation for multi-session mapping. Notably both recall *and* precision increase with more sessions. Further, we computed reconstructions using different submap sizes of $S = 1, 5, 10, 20$ keyframes each. Larger submaps are more compact as the relative overlap between consecutive maps is smaller, resulting in less data transmission and storage, while shorter submaps may be




Fig. 4: Reconstruction results for varying number of sessions and submap sizes on HD-Map dataset (a) Increasing the number of sessions improves not just semantic F-Score, but also recall *and* precision. (b) Semantic F-Score for varying submap sizes leading to different amount of transmitted data per mapped area.




Fig. 5: (a) Comparison of semantic F-Score for different HashGrid codebook sizes N_C and OcTree feature sizes N_F resulting in different map sizes (neural field weights).

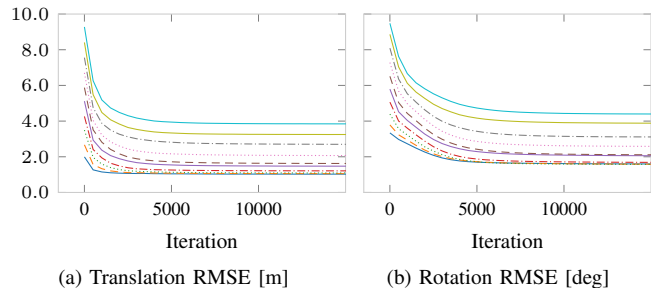


Fig. 6: Comparison of (a) translational and (b) rotational alignment RMSE over iterations for different initial offsets.

less prone to accumulating drift and fusion errors. We can observe in Fig. 4b that larger submaps perform just as good or better than smaller submaps, while requiring significantly less data per kilometer of mapped road. However, the data overhead advantage becomes less relevant for sizes $S > 10$. Interestingly, larger submaps lead to slightly superior reconstruction results compared to smaller ones, which we attribute to better alignment due to larger overlap between distinct submaps.

4) *Neural Field Grid*: In Fig. 5 we compare the effect of different HashGrid and OcTree parameterizations on the




Fig. 7: (a) Visualization of semantic (top left) and confidence score prediction (bottom right) with warmer colors denoting higher confidence. (b) Semantic scores for varying confidence thresholds c_{th} .

reconstruction quality and the size of the neural field in terms of storage. We can clearly see that HashGrids are much more efficient in our setup compared to OcTrees, requiring around $5\times$ less storage to achieve the same reconstruction quality.

5) *Alignment Robustness*: Since our joint reconstruction and pose-alignment is optimized towards a local minimum, reasonable initialization of poses is crucial to achieve good alignment. We investigate the effect of inaccurate initial poses for various amounts of disturbance on KITTI GPS poses in Fig. 6. We apply varying amount of artificial noise $e_{trans} \sim \mathcal{N}(0, \sigma_{trans})$ and $e_{rot} \sim \mathcal{N}(0, \sigma_{rot})$ for translational and rotational components with $\sigma_{trans} \in [1, 5.5]m$, $\sigma_{rot} \in [1, 5.5]^\circ$. While low to medium levels of disturbance all convergence well to similar low final alignment errors, we find that the alignment result degrades well for increasing initial errors. Even for extreme levels of initial disturbance, there is a significant reduction in pose error, proving the robustness of joint reconstruction and alignment using neural field fusion.

6) *Confidence*: We show a visualization of our confidence prediction in Fig. 7 combined with semantic prediction. We observe that the main road areas carry a high confidence, visualized in red, while areas that are rarely observed or occluded have a low-confidence prediction shown in blue. Fig. 7b depicts the reconstruction F-score, precision and recall for various thresholds $c_{th} \in [0, 1]$. We note that higher thresholds increase precision (accuracy) while decreasing recall (completeness) and find $c_{th} = 0.7$ as balance that achieves good results across our datasets.

V. CONCLUSION

In this paper, we presented a novel multi-session mapping approach for crowd-sourced map learning. Our system fuses submaps reconstructed by individual vehicles into a coherent dense semantic map that exceeds the fidelity of existing approaches for large-scale map learning. Using meshes for transmitting map data to the mapping server enables lightweight communication and re-uses locally fused information from arbitrary sensors and reconstruction methods. In turn, our neural field fusion facilitates large-scale joint alignment and 3D reconstruction without the need for dedicated

localization landmarks. In our experiments, we demonstrated the quality and robustness of the fused maps and showed the practicality and effectiveness of our approach for high-fidelity map learning. We see neural implicit representations as exchange format from vehicle to server as well as change detection as promising directions for future research.

REFERENCES

- [1] O. Dabeer, W. Ding, R. Gowaiker, S. K. Grzechnik, M. J. Lakshman, S. Lee, G. Reitmayr, A. Sharma, K. Somasundaram, R. T. Sukhvasi, and X. Wu, "An end-to-end system for crowdsourced 3d maps for autonomous vehicles: The mapping component," *2017 IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, 2017.
- [2] D. Pannen, M. Liebner, and W. Burgard, "Lane marking learning based on crowdsourced data," *2019 IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, 2019.
- [3] M. Herb, T. Weiherer, N. Navab, and F. Tombari, "Crowd-sourced semantic edge mapping for autonomous vehicles," *2019 IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, 2019.
- [4] K. Kim, S. Cho, and W. Chung, "Hd map update for autonomous driving with crowdsourced data," *IEEE Robotics and Automation Letters*, 2021.
- [5] J. J. Park, P. Florence, J. Straub, R. Newcombe, and S. Lovegrove, "DeepSDF: Learning continuous signed distance functions for shape representation," *2019 IEEE/CVF Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [6] T. Takikawa, J. Litalien, K. Yin, K. Kreis, C. Loop, D. Nowrouzezahrai, A. Jacobson, M. McGuire, and S. Fidler, "Neural geometric level of detail: Real-time rendering with implicit 3d shapes," *2021 IEEE/CVF Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2021.
- [7] T. Müller, A. Evans, C. Schied, and A. Keller, "Instant neural graphics primitives with a multiresolution hash encoding," *ACM Transactions on Graphics*, 2022.
- [8] T. Qin, Y. Zheng, T. Chen, Y. Chen, and Q. Su, "A light-weight semantic map for visual localization towards autonomous driving," *2021 IEEE Intl. Conf. on Robotics and Automation (ICRA)*, 2021.
- [9] D. Pannen, M. Liebner, W. Hempel, and W. Burgard, "How to keep hd maps for automated driving up to date," *2020 IEEE Intl. Conf. on Robotics and Automation (ICRA)*, 2020.
- [10] Q. Cheng, N. Zeller, and D. Cremers, "Vision-based large-scale 3d semantic mapping for autonomous driving applications," *2022 Intl. Conf. on Robotics and Automation (ICRA)*, 2022.
- [11] J. L. Schonberger and J.-M. Frahm, "Structure-from-motion revisited," *2016 IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [12] K. Wang, F. Gao, and S. Shen, "Real-time scalable dense surfel mapping," *2019 Intl. Conf. on Robotics and Automation (ICRA)*, 2019.
- [13] M. Herb, T. Weiherer, N. Navab, and F. Tombari, "Lightweight semantic mesh mapping for autonomous vehicles," *2021 IEEE Intl. Conf. on Robotics and Automation (ICRA)*, 2021.
- [14] H. Oleynikova, Z. Taylor, M. Fehr, R. Siegwart, and J. Nieto, "Voxblox: Incremental 3d euclidean signed distance fields for on-board mav planning," *2017 IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, 2017.
- [15] V. Reijgwart, A. Millane, H. Oleynikova, R. Siegwart, C. Cadena, and J. Nieto, "Voxgraph: Globally consistent, volumetric mapping using signed distance function submaps," *IEEE Robotics and Automation Letters*, 2020.
- [16] A. Rosinol, M. Abate, Y. Chang, and L. Carlone, "Kimera: an open-source library for real-time metric-semantic localization and mapping," *2020 IEEE Intl. Conf. on Robotics and Automation (ICRA)*, 2020.
- [17] H. Hu, H. Yang, J. Wu, X. Lei, F. Bieder, J.-H. Pauls, and C. Stiller, "Large-scale 3d semantic reconstruction for automated driving vehicles with adaptive truncated signed distance function," *2023 IEEE Intelligent Vehicles Symposium (IV)*, 2023.
- [18] B. Mildenhall, P. P. Srinivasan, M. Tancik, J. T. Barron, R. Ramamoorthi, and R. Ng, "Nerf: Representing scenes as neural radiance fields for view synthesis," *European Conf. on Computer Vision (ECCV)*, 2020.
- [19] M. Tancik, V. Casser, X. Yan, S. Pradhan, B. P. Mildenhall, P. Srinivasan, J. T. Barron, and H. Kretschmar, "Block-nerf: Scalable large scene neural view synthesis," *2022 IEEE/CVF Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2022.
- [20] K. Rematas, A. Liu, P. Srinivasan, J. Barron, A. Tagliasacchi, T. Funkhouser, and V. Ferrari, "Urban radiance fields," *2022 IEEE/CVF Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2022.
- [21] X. Zhong, Y. Pan, J. Behley, and C. Stachniss, "Shine-mapping: Large-scale 3d mapping using sparse hierarchical implicit neural representations," *2023 IEEE Intl. Conf. on Robotics and Automation (ICRA)*, 2023.
- [22] C.-H. Lin, W.-C. Ma, A. Torralba, and S. Lucey, "Barf: Bundle-adjusting neural radiance fields," *2021 IEEE/CVF Intl. Conf. on Computer Vision (ICCV)*, 2021.
- [23] J. Deng, Q. Wu, X. Chen, S. Xia, Z. Sun, G. Liu, W. Yu, and L. Pei, "Nerf-loam: Neural implicit representation for large-scale incremental lidar odometry and mapping," *2023 IEEE/CVF Intl. Conf. on Computer Vision (ICCV)*, 2023.
- [24] R. Elvira, J. D. Tardós, J. M. M. Montiel, J. D. Tardos, and J. M. M. Montiel, "Orbslam-atlas: a robust and accurate multi-map system," *2019 IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems*, 2019.
- [25] P. Schmuck and M. Chli, "Ccm-slam: Robust and efficient centralized collaborative monocular simultaneous localization and mapping for robotic teams," *Journal of Field Robotics*, 2019.
- [26] R. Mur-Artal, J. M. M. Montiel, and J. D. Tardos, "Orb-slam: A versatile and accurate monocular slam system," *IEEE Transactions on Robotics*, 2015.
- [27] A. Cramariuc, L. Bernreiter, F. Tschoop, M. Fehr, V. Reijgwart, J. Nieto, R. Siegwart, and C. Cadena, "maplab 2.0 – a modular and multi-modal mapping framework," *IEEE Robotics and Automation Letters*, 2023.
- [28] Y. Chang, Y. Tian, J. P. How, and L. Carlone, "Kimera-multi: a system for distributed multi-robot metric-semantic simultaneous localization and mapping," *2021 IEEE Intl. Conf. on Robotics and Automation (ICRA)*, 2021.
- [29] S. Golodetz, T. Cavallari, N. A. Lord, V. A. Prisacariu, D. W. Murray, and P. H. S. Torr, "Collaborative large-scale dense 3d reconstruction with online inter-agent pose optimisation," *IEEE Transactions on Visualization and Computer Graphics*, 2018.
- [30] R. Dubois, A. Eudes, J. Moras, and V. Fremont, "Dense decentralized multi-robot slam based on locally consistent tsdf submaps," *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, 2020.
- [31] X. Liu, W. Ye, C. Tian, Z. Cui, H. Bao, and G. Zhang, "Coxgraph: Multi-robot collaborative, globally consistent, online dense reconstruction system," *2021 IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, 2021.
- [32] R. Mur-Artal and J. D. Tardós, "Orb-slam2: An open-source slam system for monocular, stereo, and rgb-d cameras," *IEEE Transactions on Robotics*, 2017.
- [33] A.-Q. Cao and R. de Charette, "Monoscene: Monocular 3d semantic scene completion," *2022 IEEE/CVF Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2022.
- [34] Z. Li, T. Müller, A. Evans, R. H. Taylor, M. Unberath, M.-Y. Liu, and C.-H. Lin, "Neuralangelo: High-fidelity neural surface reconstruction," *2023 IEEE/CVF Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2023.
- [35] A. Gropp, L. Yariv, N. Haim, M. Atzmon, and Y. Lipman, "Implicit geometric regularization for learning shapes," *Proceedings of the 37th Intl. Conf. on Machine Learning*, 2020.
- [36] T. Takikawa, O. Perel, C. F. Tsang, C. Loop, J. Litalien, J. Tremblay, S. Fidler, and M. Shugrina, "Kaolin wisp: A pytorch library and engine for neural fields research," <https://github.com/NVIDIAGameWorks/kaolin-wisp>, 2022.
- [37] Y. Liao, J. Xie, and A. Geiger, "Kitti-360: A novel dataset and benchmarks for urban scene understanding in 2d and 3d," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022.
- [38] P. Besl and N. D. McKay, "A method for registration of 3-d shapes," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 1992.
- [39] A. V. Segal, D. Haehnel, and S. Thrun, "Generalized-icp," *Robotics: Science and Systems*, 2010.
- [40] M. Kazhdan and H. Hoppe, "Screened poisson surface reconstruction," *ACM Transactions on Graphics*, 2013.
- [41] A. Knapitsch, J. Park, Q.-Y. Zhou, and V. Koltun, "Tanks and temples," *ACM Transactions on Graphics*, 2017.