Autonomous Navigation in Inclement Weather Based on a Localizing Ground Penetrating Radar

Teddy Ort, Igor Gilitschenski, and Daniela Rus

***Abstract*—Most autonomous driving solutions require some method of localization within their environment. Typically, onboard sensors are used to localize the vehicle precisely in a previously recorded map. However, these solutions are sensitive to ambient lighting conditions such as darkness and inclement weather. Additionally, the maps can become outdated in a rapidly changing environment and require continuous updating. While LiDAR systems don't require visible light, they are sensitive to weather such as fog or snow, which can interfere with localization. In this letter, we utilize a Ground Penetrating Radar (GPR) to obtain precise vehicle localization. By mapping and localizing using features beneath the ground, we obtain features that are both stable over time, and maintain their appearance during changing ambient weather and lighting conditions. We incorporate this solution into a full-scale autonomous vehicle and evaluate the performance on over 17 km of testing data in a variety of challenging weather conditions. We find that this novel sensing modality is capable of providing precise localization for autonomous navigation without using cameras or LiDAR sensors.**

***Index Terms*—Autonomous vehicle navigation, field robots, wheeled robots, intelligent transportation systems.**

Fig. 1. Our LGPR based autonomous navigation system is capable of successfully localizing from a single map in different challenging weather and lighting conditions including snow, rain, and darkness.

I. Introduction

ROBUST localization in diverse conditions is a key challenge to enable the widespread deployment of autonomous vehicles. Since relying purely on a global navigation satellite system (GNSS), such as the Global Positioning System (GPS) does not provide sufficient precision, research and industry efforts have focused primarily on utilizing cameras and laser scanners for the navigation task. These systems typically use Simultaneous Localization and Mapping (SLAM) algorithms [1] for creating and maintaining maps of the environment that allow for highly precise localization and navigation.

Using vision and laser sensors as localization sources comes with its own set of challenges. For instance, localization from purely visual maps needs to account for the fact that, even in the absence of occlusions, the appearance of the environment strongly varies depending on weather, season, time of day, and a potentially changing environment [2]. While laser scanners do not suffer from changing illumination conditions, they are still highly affected by weather, occlusions, and environment dynamics.

Recently, the use of a Ground Penetrating Radar (GPR) was proposed for localization [3]. Similar to the aforementioned approaches, this Localizing Ground Penetrating Radar (LGPR) can be combined with GPS to constrain the search space. The main advantages of using GPRs for the localization task is their robustness to dynamic surface environments, illumination changes, and visibility conditions. This is made possible by a ground facing sensor array that creates maps of the underlying structure of the soil and uses these maps for subsequent localization. However, the LGPR system has never been evaluated on an autonomous vehicle in adverse weather conditions which include some of the key applications for this technology.

In this work, we deploy the LGPR on a full-scale autonomous Toyota Prius (as shown in Fig. 1). We describe its incorporation into a novel autonomous driving pipeline and justify all of our design choices. Through extensive evaluations, we demonstrate the practical capability of LGPR to enable an autonomous vehicle to drive without relying on any visual or surface features. For this reason, the method works even when the driving surface is covered with rain or snow. While the LGPR performs vehicle localization, it does not enable dynamic obstacle detection in inclement weather. However, in this letter we focus on the localization problem and leave the detection of dynamic obstacles in challenging weather conditions for future work.

Manuscript received September 10, 2019; accepted February 7, 2020. Date of publication February 26, 2020; date of current version March 9, 2020. This letter was recommended for publication by Associate Editor P. Tokekar and Editor D. Popa upon evaluation of the reviewers' comments. This work was supported by MIT Lincoln Laboratory. *(Corresponding author: Teddy Ort.)*

The authors are with the Computer Science and Artificial Intelligence Lab, Massachusetts Institute of Technology, Cambridge, MA 02139 USA (e-mail: teddy@mit.edu; igilitschenski@mit.edu; rus@mit.edu).

This letter has supplementary downloadable material available at http://ieeexplore.ieee.org, provided by the authors.

Digital Object Identifier 10.1109/LRA.2020.2976310

Overall, our contributions can be summarized as follows:

- Design of a new autonomous driving pipeline based on LGPR.
- Demonstrate mapping, localization, planning, and control of an autonomous vehicle without relying on any visible surface features.
- Extensive evaluation of our system localizing on the same map in a variety of weather conditions including clear weather, rain, and snow.

II. Related Work

Radar-Based Perception. Radars serve as a key component for advanced driver assistance systems and autonomous vehicles [4]. Past research in that space focused on fusing radar data with various other sensor types particularly Lidars [5] and cameras [6] for applications such as pedestrian detection [7] and tracking [8]. Radar-based SLAM approaches [9] usually model structural features in the surrounding environment and have also been investigated in the context of vehicle self-localization [10], [11]. However, these approaches may fail without special considerations to account for dynamic objects and changing surroundings.

Appearance and Persistence Modeling. Particularly for visual SLAM [1], [2], an extensive body of research explicitly focuses on addressing these types of problems. The breadth of approaches involves inpainting and removal of dynamic objects [12], [13], selecting particularly persistent landmarks for map storage [14], and selecting appearance specific landmarks [15] or map segments [16]. We avoid the problems addressed in this line of work altogether by using sub-surface features that do not suffer from frequent appearance changes and occlusions.

GPRs in Robotics. Sub-surface sensing tasks are usually carried out using GPRs. So far, corresponding research work mainly focused on applications such as measuring the makeup and content of the soil [17], [18], locating underground structures [19], mapping archaeological features [20], or landmine detection [21], [22]. Similarly, GPR deployment on robots mainly focused on data collection tasks such as autonomous surveys [23], [24], inspection [25], or lunar crust analysis [26]. Contrary to these approaches, we do not focus on classifying the content in the GPR images but rather on using them to improve the navigation capabilities of the autonomous system.

Only a limited number of works focuses on consideration of occluded structures in radar-based SLAM such as [27]. The first use of a GPR for localization was reported in [3] focusing on describing the design of a low-profile low power GPR system and first demonstrations of real-time localization. In contrast, the present work demonstrates a fully-integrated LGPR system for autonomous driving. Moreover, we establish and evaluate the sensor's suitability for successful localization across different weather conditions even when using the same map.

III. System Overview

An autonomous vehicle requires a large system architecture to enable safe navigation. The key components of this system include: mapping, localization, planning, and control. Fig. 2 shows a diagram of the various system components. In the following subsections, we describe the function of these components with an emphasis on how the design differs in an LGPR system.

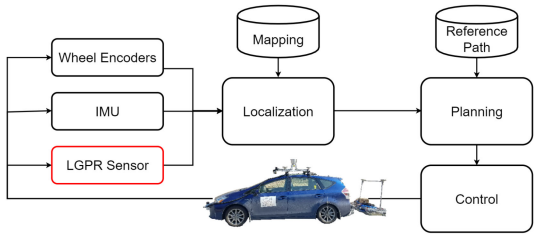

Fig. 2. The autonomous navigation system architecture showing the key system components: Mapping, Localization, Planning, and Control. The LGPR system used as the primary localization sensor is outlined in red.

A. Mapping

During the mapping phase, the vehicle is driven by a human operator and the data from the vehicle's sensors is recorded. Once this step is completed, the saved map enables the vehicle to localize in that area. There are two key differences between the dense 3D maps required in typical vision or LiDAR systems, and those created using an LGPR sensor. Firstly, because vision and LiDAR maps record surface features, they require maintenance when any of those features change which can result in burdensome repeated map updates. LGPR maps on the contrary, record subsurface features which are unlikely to change frequently. Secondly, dense sensor maps typically take up large amounts of space which makes them difficult to store and transmit. For example, [28] compares the size of a typical (20,000 mile) topological map such as OpenStreetMap with that of a highly compressed 2D sensor map of the same area. The topological map only requires 3 GB while the sensor map requires about 200 GB of storage space. However, the topological map does not contain enough information for precise localization. To store an LGPR map of the same area we require approximately 160 GB. The data structure used to store the LGPR maps is described in Sec IV-B2. Intuitively, these are smaller because the sensor measures only a thin slice directly below the vehicle, while typical 3D maps contain a detailed view of the entire environment including surrounding buildings and vegetation. Thus, LGPR maps can provide precise localization in changing surface conditions without requiring as much storage space.

B. Localization

During operation, the autonomous vehicle must have precise localization information in order to ensure that it closely follows the desired path. Standard freeway lane-widths in the United States are 3.6 m [29] while the width of a typical passenger vehicle is 1.7 m. This gives only about 1m of clearance on each side which is beyond the accuracy of GPS systems even when augmented with wheel odometers [30]. LiDAR or vision based maps provide localization with sufficient accuracy. However, they are susceptible to ambient lighting conditions and fail if road surface features such as lane markings are obscured by

rain or snow. The LGPR sensor provides precise localization to the vehicle by matching the sensor data during operation to that stored during the mapping phase. Finally, this localization update is then probabilistically fused with the onboard proprioceptive sensors including the wheel odometry and the Inertial Measurement Unit (IMU). These are combined to provide a robust and precise localization estimate at a rate of 40 Hz. Note that distance travelled during a single cycle is given by $d = v/f$ where v is the vehicle velocity and f is the loop frequency. Therefore, $f = 40$ Hz is enough to ensure that even at highway speeds of 30 m/s the vehicle would travel <1 m before receiving a new localization estimate.

C. Planning

The planning component determines the path the vehicle is required to take to safely navigate to its destination. It begins by determining the current location of the vehicle which it receives from the localization system see Sec III-B. Next, it uses the LGPR-map from Sec III-A to determine which roads must be traversed, and dynamic obstacle detection systems to plan a collision free path toward the goal. Choosing which roads to take is performed using a shortest-path algorithm such as Dijkstra [31] and is similar to that performed on standard vehicle navigation devices thus we will not expound on it here. However, autonomously following the desired path is not trivial. To do so, we utilize a pure pursuit controller see Sec. V-A for details. This algorithm is capable of planning paths that follow a reference trajectory even when the vehicle cannot exactly follow the reference as is typical in real-world scenarios. This algorithm is a challenging test for the LGPR localization system, which in previous work [3] was only evaluated while being manually driven exactly over the same path as that in the map.

D. Control

The control component is required to determine the actual control values to transmit to the vehicle, including accelerator, brake, and steering wheel angle. These values are computed in order to ensure that the vehicle will follow the planned trajectory. In our system, we utilize two PID controllers to obtain the necessary control values. One of these controls the accelerator and the brake, while the other controls the steering. The gains were tuned to allow the vehicle to quickly achieve a desired configuration while avoiding any instabilities due to lag in the system.

IV. A Localizing Ground Penetrating Radar

In this section, we describe how the LGPR sensor is used both to build maps during the mapping phase, as well as to provide localization during autonomous operation (see Fig. 4). The LGPR sensor uses electromagnetic radiation in the 100–400 MHz range which is much lower than typical (1–3 GHz) GPR devices used for surveying [32]. This allows the sensor to resolve subsurface geological features in the 20–30 cm range with the additional benefit of improving the penetration depth. The shallow penetration depth of higher frequency sensors would not penetrate below the road as necessary for observing unique underground features. In general, the penetration depth is dependent on the relative permeability of the soil [3]. In our testing region, (Devens, Massachusetts) skin depths of $D \approx 100$ cm are typical leading to penetration depths in the 2–3 m range. Road materials and soil are transparent to radar signals at these frequencies. However, changes in the dielectric properties of the subsurface geology due to variations in soil type and density, rocks, roots, or layer boundaries cause the signals to be reflected back to the LGPR sensor. During each element pulse, the neighboring element is used to measure the intensity of the reflected signal and record an "image" of the subterranean content beneath the element at the time the pulse occurred. In the following subsections we describe the key features of A) the LGPR sensor utilized, and B) the localization algorithm used both to record the prior map, as well as to localize a new scan.

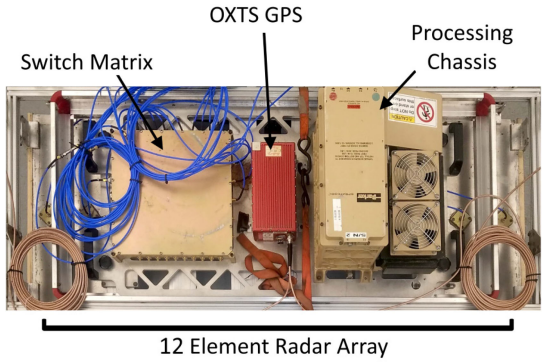

Fig. 3. The LGPR sensor mounted on the rear of the autonomous vehicle. The main components include the 11 element radar array and the switch matrix which collected the signals from the array. The processing chassis contains the onboard computer for the sensor data processing, and the GPS unit used for labeling the prior maps and initializing the search window for localization.

A. LGPR Sensor

The LGPR sensor used in this work was custom designed by MIT Lincoln Laboratory for the purpose of localization. Fig. 3 shows the key sensor hardware components. The sensor measures 152 cm $\times$ 61 cm $\times$ 7.6 cm which requires it to be mounted outside the vehicle. Because the measurement range resolution depends on the width of the array, the sensor cannot be easily miniaturized. Therefore, the sensor was mounted on the rear of the test vehicle at 32 cm above the ground (See Fig. 1). This distance was chosen because it is close to the ground - ensuring a greater penetration depth - while still remaining high enough for adequate road clearance. Note that in [3] the sensor was able to be lowered to 15.24 cm which may have an effect on our results. Furthermore, we believe the placement on the rear of the vehicle is most stable since the vehicle begins to depart from the recorded path front first. This ensures the localization estimates will be available as long as possible to correct for any departures. However, due to hardware limitations, the sensor location remained fixed so we cannot know empirically how strong these effects are.

Here we will describe the key features of the sensor, see [3] for more detailed specifications. The LGPR array is composed of 12 radar elements. During a single sensor sweep, each element

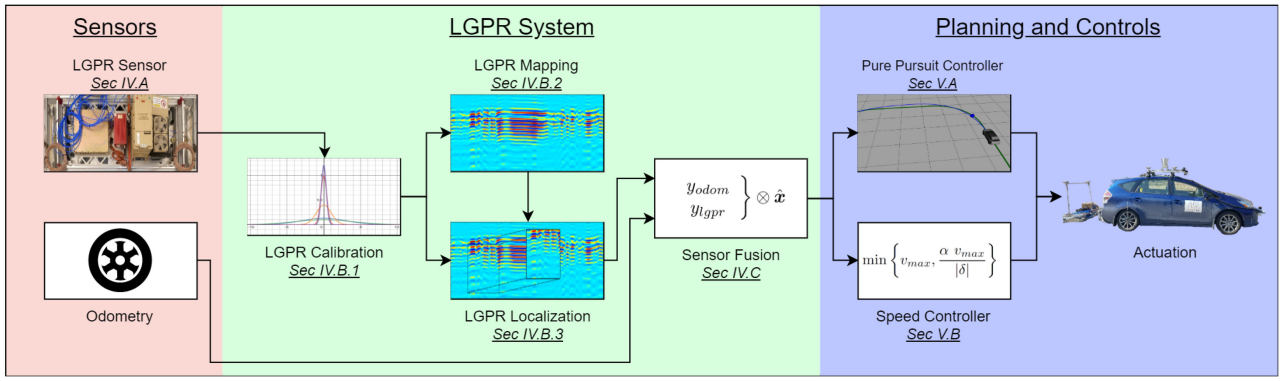

Fig. 4. The autonomous navigation pipeline which processed sensor inputs from the wheel odometry, and LGPR sensor, and computed steering and speed commands to send to the vehicle to autonomously drive the vehicle along the goal path.

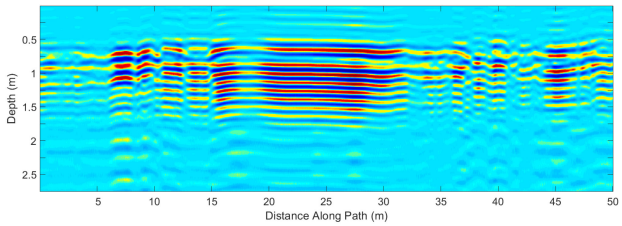

Fig. 5. LGPR data collected from a single radar element as the vehicle was driven over a path. The colors indicate the intensity of the reflected signal at each depth and the high degree of structure in the data allows the subterranean geological features to be coherently measured using this method.

transmits for a fixed period while its neighboring element receives the reflected signal. Thus, although there are 12 elements only 11 channels of data are returned. The data is then passed through a number of post-processing steps including an Infinite Impulse Response filter, online calibration, as well as factory measured calibration corrections to ensure the data is free of noise and sensor biases. The resulting measurement from each channel is an array of 369 bytes representing the relative intensities of the reflections at each depth. Each channel pulses for 8 ms which gives a maximum sensor rate of 126 Hz. Thus, a complete sweep of the sensor can be considered an image of single-byte pixels with dimension 11×369 where each pixel color gives the intensity of the reflected wave at a particular location. Fig. 5 shows an example of the values of a single element taken as the vehicle moved along a path. The coherent structures in these images provide a picture of the content of the subterranean environment measured by the sensor.

B. Localization Algorithm

The onboard computer in the processing chassis receives the sensor data from the switch matrix, which collects the signals from each of the 12 radar elements. The sensor has three modes of operation: 1) Calibration, 2) Map Creation, and 3) Localization.

1) Calibration: During sensor calibration, each element is pulsed for a short period and the mean values are recorded to ensure the resulting signals are mean centered at 0 and to compensate for variations in temperature. The calibration routine is fully automated, only takes a few seconds, and [3] found that it allowed the sensor to operate robustly at temperatures ranging from $-5\,^{\circ}$C to $50\,^{\circ}$C. In our experiments, we recalibrate daily, which is necessary to compensate for thermal expansion of the sensor components. However, since the recalibration always zeros the mean of the signal, we found it does not preclude localizing to a prior map created with a prior calibration.

2) Map Creation: During the map creation phase, a human driver operates the vehicle and each scan from the sensor is recorded in sequence building a 3D tunnel of dimension $11 \times 369 \times N$ for the 11 channels, 369 discrete depth bins, and N sensor sweeps. Additionally, the GPS coordinates of the vehicle are collected and associated with each scan. It is important to notice that a standard GPS device could be used for this labeling process, and in fact, the LGPR localization would still outperform even the GPS device that was used in the making of the map itself. Consider that the key information the LGPR sensor must provide during localization is the pose of the vehicle with respect to the original data. There is no real requirement to use GPS coordinates at all when labeling the scans. While it is convenient to use GPS coordinates in order to allow the sensor localization results to have some real-world meaning, even if the mapping GPS is not accurate to the "true" GPS coordinates, the vehicle will still be able to follow the desired path. For this reason, the extrinsic calibration between the GPS and LGPR is also not needed. Since the GPS is only being used for labeling, a fixed offset will not affect system performance. We will call these coordinates LGPR-GPS coordinates to differentiate them from the true GPS coordinates that are fixed to the earth frame.

3) Localization: During the localization phase, a single sweep of the sensor is compared to the prior map database. In order to keep memory costs constant, the map data is pulled from the database in square grids in the local area where the vehicle is operating. Next, for a given hypothetical pose for the current scan in the map, the correlation is calculated as

$$r_{A,B} = \frac{\sum_{i,d} A_{i,d} B_{i,d}}{\sqrt{\sum_{i,d} A_{i,d}^2 B_{i,d}^2}} \tag{1}$$

where A, and B represent the current scan, and the scan data from the prior map respectively, i spans the number of channels, and d spans the number of depth bins. Note, that rather than assuming the vehicle sensor will directly overlap a previous scan, the scan from the prior map is generated for any 5 degrees of freedom (DOF) pose (x, y, roll, pitch, yaw) of the vehicle by interpolating in the 3D tunnel from the dataset. (Even though the vehicle does change altitude, since it's height with respect to the map remains constant the z value is unnecessary. However, the roll and pitch are included because small variations in the orientation of the sensor due to vehicle suspension could lead to large errors in the position of underground features.) In this way, any hypothetical pose that overlaps the data is considered. To find the optimal match that maximizes $r_{A,B}$ in (1), a particle swarm optimization is performed to balance the need for quickly finding solutions when they exist in the expected location, with the ability to search over a large area. As described in [3] the number of particles, the size of the search region, and the number of iterations, are tuned online to ensure fast operation when the correlations are high, with a more exhaustive search if the correlations indicate a poor match. Importantly, although the sensor itself, polls at a constant 126 Hz, the particle swarm optimization can limit the actual localization frequency. With the computational resources available (see Sec. VI-A), we found that the localization results could keep up with the sensor frequency only when the vehicle was following closely to the path in the prior map. However, when correlations were lower due to noise, or the sensor not fully overlapping with the prior map, the localization frequency would drop.

Each localization solution includes a 5 DOF pose composed of LGPR-GPS coordinates and orientation, the correlation value $r_{A,B} \in [-1, 1]$, and an overlap value $u_{A,B} \in [1..11]$ indicating how many of the 11 channels overlapped with the prior map in the optimal pose. In the next section, we detail how these localization estimates were integrated into the autonomy pipeline.

C. Sensor Fusion

We utilize multiple instances of the Extended Kalman Filter (EKF) algorithm for fusing the various sensors. One instance fuses the wheel encoders and IMU sensor measurements to obtain a filtered odometry estimate, and a second instance fuses these with the LGPR measurements to obtain an estimate in the LGPR-GPS reference frame. The motivation for using two separate EKF instances is twofold. Firstly, the proprioceptive sensors run at a consistently fast frequency while the LGPR sensor is sometimes slower if the correlation quality is low. Secondly, this gives access to a smooth position estimate in the local frame of the robot, while the global estimate can have discrete jumps if an LGPR correction comes in after a relatively long pause. The EKF algorithm has been exhaustively explored in the literature [33]–[35] so we will only cover the key parts that are important for our implementation. We utilize the EKF architecture as described in [36] for both instances. The state vector is 8 dimensional containing the 3 DOF pose of the robot, its velocities, and linear accelerations:

$$\boldsymbol{x} = \begin{bmatrix} x, & y, & \theta, & \dot{x}, & \dot{y}, & \dot{\theta}, & \ddot{x}, & \ddot{y} \end{bmatrix}^\top$$

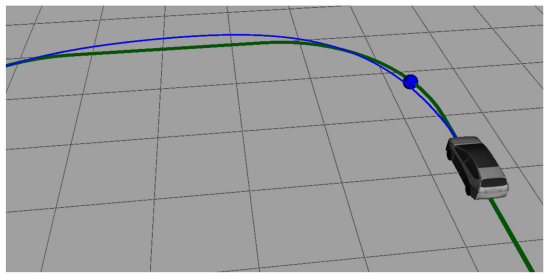

Fig. 6. The pure pursuit tracking algorithm was used to smoothly steer the vehicle along the goal path (shown in green). The blue sphere shows the lookahead distance and the blue curve shows the path the controller intends to take. As expected, this path tends to deviate from the goal path in turns.

We utilize the wheel encoders only to measure the linear velocity $\dot{x}$ and the IMU to measure θ and $\dot{\theta}$. Note the IMU is paired with a magnetometer giving it an absolute measurement of θ rather than solely a relative one. We extrinsically calibrate the IMU using the procedure described in [37]. Each LGPR measurement contains

$$\boldsymbol{z} = \begin{bmatrix} \text{latitude}, & \text{longitude}, & \theta \end{bmatrix}^\top$$

which is transformed into the Universal Transverse Mercator (UTM) frame using the transform in [38]. The process noise covariance matrix $\boldsymbol{w} \in \mathbb{R}^{8\times 8}$ was tuned in the field to match observed uncertainties. It contains only diagonal entries with the values:

$$\text{diag}(\boldsymbol{w}) = [0.5, \quad 0.5, \quad 6, \quad 2.5, \quad 2.5, \quad 2, \quad 1, \quad 1] \times 10^{-2}$$

Finally, since each of the sensors directly measures a state variable, the associated observation matrix H (the jacobian of the observation function) is all 0's except for 1's in the diagonal positions of the measured state variables.

V. Planning and Control Pipeline

A. Pure Pursuit Controller

The pure pursuit controller [39] is a path tracking controller for Ackermann steered autonomous vehicles. It was used by the MIT entry in the DARPA Urban Challenge in 2007 [40] and has been widely adopted for autonomous driving applications. One key feature of the algorithm is it's pursuit of a "lookahead point" on the goal path at some distance d ahead of the vehicle to smoothly adapt to deviations of the vehicle from the reference path. The pure pursuit steering equation is

$$\delta = \tan^{-1}\left(L\frac{2x}{d^2}\right)$$

Where δ is the required steering angle, x is the offset to the goal point, L is the vehicle length (measured between the front and rear axles) and d is the lookahead distance. Fig. 6 shows the vehicle driving along the goal path (in green) and approaching a left turn. The blue path shows how the pure-pursuit algorithm will steer the vehicle. As expected, it begins by oversteering in the beginning of the turn and then overshoots at the end. The amount of path deviation depends on the lookahead parameter d. However, when utilizing LGPR, deviations from the path cause

the sensor overlap value $u_{A,B}$ to drop as discussed previously. The EKF implementation described in the previous section is crucial to allow the vehicle to continue navigation even when the overlap is low.

B. Speed Controller

For the purpose of evaluating the LGPR's capability to provide localization for autonomous navigation, the speed of the vehicle is not a critical element. However, a controller to regulate the speed was still required to maintain the vehicle cruising speed and appropriately decelerate in turns. To this end, we designed a speed control law:

$$v = \min\left\{v_{max}, \frac{\alpha v_{max}}{|\delta|}\right\}$$

Where v_{max} is the maximum speed, α is a tunable parameter, δ is the steering angle, and v is the commanded speed sent to the vehicle controls. [3] evaluate the performance of the localization with a human operator at speeds of up to 65 mph. For our tests however, using a fully autonomous vehicle, we set $v_{max} = 10 - 20$ mph for safety reasons. Finally, we found that $\alpha = 0.1$ provided an appropriate rate of speed reduction while turning.

VI. Evaluation

We performed extensive evaluation of the system in order to verify that the LGPR sensor could provide a localization estimate precise enough for autonomous navigation. We measure the localization accuracy while the system is being operated autonomously and compare it to the localization accuracy when a human is driving. Finally, because the key benefit of LGPR localization over existing localization techniques is its capability in inclement weather, we compare the localization accuracy in different weather conditions.

A. Setup

The testing site was a closed rural area in Devens, Massachusetts. It includes ∼7 km of unmarked roads and 9 intersections. Testing took place over a period of six months to compare a variety of weather and lighting conditions including clear weather, snow, rain, and darkness (See Fig. 1).

The LGPR sensor was mounted on the trailer hitch of a Toyota Prius which was modified to drive by wire. Two computers ran in parallel to enable the autonomous system. Onboard the LGPR sensor the *LGPR-PC* carried out all localization related computation. The sensor broadcast localization estimates via Ethernet to the *PRIUS-PC* onboard the vehicle. The *PRIUS-PC* then incorporated that information into the autonomous driving system. Note that the *PRIUS-PC* was a consumer grade laptop computer with an Intel Xeon E3-1505Mv5 Quad Core 2.80 GHz processor, 32 GB of RAM, and running Ubuntu 16.04. The only other sensors used during testing were two wheel encoders on each of the rear wheels and an IMU.

To evaluate the performance of the system, an additional RTK-GPS unit with Differential GPS (DGPS) corrections from a fixed base station was installed on the vehicle for ground truth. The setup gives the precise vehicle position with ∼2 cm accuracy. While the LGPR system already uses an onboard GPS for creating maps, and initializing the search radius, that device did not have access to the base station correction in order to simulate real-world operation. The base station only has a broadcast radius of up to 9 km, which renders it impractical for use in general purpose autonomous vehicles. The measurements from the base station were used only for analysis and evaluation of the LGPR system.

We chose to use the uncorrected GPS unit for the LGPR device even during the mapping phase as requiring an extremely accurate GPS sensor could make it very difficult to create maps. Furthermore, since the GPS coordinates saved are only necessary to be locally accurate for navigation purposes, it is not necessary that the LGPR-GPS coordinates match the true GPS coordinates.

B. Testing

During testing, each test was composed of two runs each of which yielded two trajectory measurements for a total of four trajectories per test. The first run was a mapping run with the LGPR system in map creation mode and the vehicle driven by a human driver. The trajectory measured by the LGPR sensor during map creation will be denoted T^{map}_{LGPR}. Additionally, the RTK-GNSS system was used to precisely track the true trajectory of the vehicle during map creation and that trajectory will be denoted T^{map}_{GNSS}. Next, the vehicle was driven over the same track again, with the LGPR sensor set in tracking mode. The localization estimates during the testing run will be denoted T^{test}_{LGPR}. Again, the RTK-GNSS was used to measure the true trajectory of the vehicle and that trajectory will be denoted T^{test}_{GNSS}. Note that during the testing run the vehicle could be piloted by either a human driver, or the autonomous system.

The LGPR sensor only needs to localize the vehicle within the coordinates of the created map, rather than in the fixed ground-truth GPS frame. Therefore, we cannot simply compare the T^{test}_{LGPR} estimate to the T^{test}_{GNSS} because the coordinates of the LGPR system are limited by the accuracy of the onboard GPS unit used for mapping (which does not receive base station corrections). Hence, there is no expectation for the T^{test}_{LGPR} values to match the "correct" ones, rather they need to be consistent with those measured in the mapping phase in order to allow accurate path following. Consequently, as in [3], we utilize a relative metric to obtain the accuracy of the LGPR localization by computing the mean relative difference. That is, the mean error over a run with n scans is computed as:

$$\frac{1}{n}\sum_{n}\left\|\left(T^{test}_{GNSS,i} - T^{map}_{GNSS,i}\right) - \left(T^{test}_{LGPR,i} - T^{map}_{LGPR,i}\right)\right\|$$

where, the ith estimate of the T^{map}_{GNSS} is denoted $T^{map}_{GNSS,i}$ etc. Note that the subtractions shown are the vector differences between the offset from the map to the test runs as measured by the LGPR sensor, and the same vector as measured by the RTK-GNSS system. This method provides a measurement of how closely the relative position estimate of the LGPR system

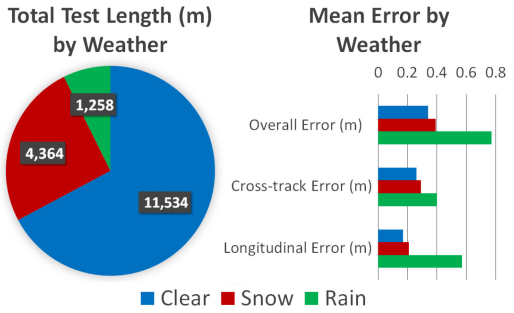

Fig. 7. A comparison of the LGPR estimates compared to ground truth in various weather conditions.

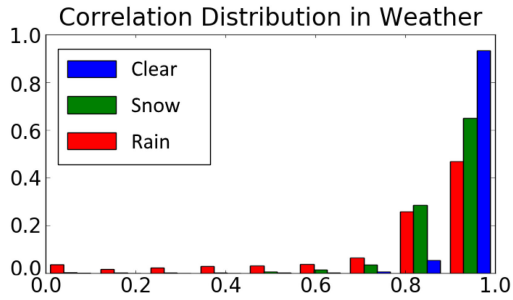

Fig. 8. A histogram showing the distribution of scan correlations between the test run and the LGPR map in various weather conditions.

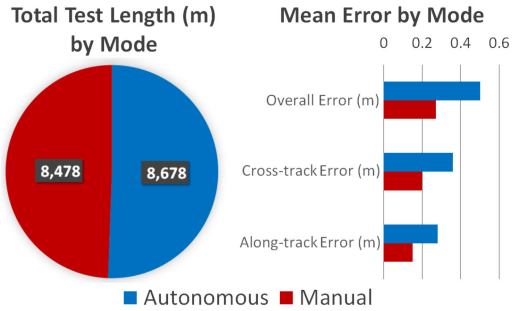

Fig. 9. A comparison of the LGPR estimates compared to ground truth in manual and autonomous driving modes.

matches the ground truth RTK-GNSS system which is necessary for autonomous path following.

One concern with calculating the mean error is aligning the timing of the trajectories. Each trajectory is measured using individual sensor clocks yielding four separate time series. The two trajectories for each run occurred at the same real world time, and therefore alignment only requires finding the scalar offset between the clocks. However, the time steps between the map and test runs cannot be as easily aligned because the runs took place over different intervals. To align these four trajectories to a unified time axis, we first align the estimates in T_{GNSS}^{map} and T_{GNSS}^{test} by aligning each point in the test trajectory with the closest point in the map. Next, the two scalar offsets that best align the LGPR and RTK-GNSS trajectories respectively were found through numerical optimization.

Finally, after synchronizing the timing and calculating the total mean error, the error vector $\mathbf{e}$ at each scan point was decomposed into a cross-track error, $\mathbf{e}_{\perp}$, and an along-track error, $\mathbf{e}_{\parallel}$ for additional insight. First, the unit tangent vector at each point was calculated using a centered finite difference along T_{GNSS}^{test} and then the components along the track and perpendicular to the track were found such that $\mathbf{e} = \mathbf{e}_{\perp} + \mathbf{e}_{\parallel}$. For autonomous navigation purposes, the requirements for low cross-track error are typically more stringent than the along-track error due to lane tolerances.

C. Results

The system was evaluated over a total of 17 km composed of both manual and autonomous driving modes in clear weather, rain, and snow.

1) Weather Conditions: Fig. 7 shows the results from comparing the effect of weather on localization accuracy. In clear weather, the system achieved 0.34 m mean total error with a cross-track error of only 0.26 m. This value is significantly better than GPS error which is typically >1 m even when fused with inertial and wheel odometry sensors [30]. Comparing the results from clear weather and snow, there is very little degradation in localization accuracy with mean total error of 0.39 m and mean cross-track error of 0.29m respectively.

In the tests during rain, a degradation in the localization quality was observed, The mean total error jumped to 0.77 m and mean cross-track error was 0.40 m. However, these values are still significantly lower than utilizing a GPS/INS solution and were not large enough to require disengagement of the autonomous system. This degradation in localization accuracy can potentially be attributed to the water in the ground changing the dielectric properties of the soil and inhibiting the penetration depth of the signal. Fig. 8 shows a histogram of the correlation values of the localization measurements as calculated in Eq. (1) in various weather conditions. It is interesting that rather than the rain causing a high instance of very low correlations, it only shifts the histogram left which indicates that the rain is not fully obscuring the signal but rather "blurring" it and making the measured scans more difficult to match to the map. It is possible that these signals could be cleaned in post-processing to achieve accuracy in rain closer to that in clear weather. However, we leave that for future work.

2) Driving Mode: Finally, Fig. 9 shows the evaluation results when comparing autonomous and manual driving. During manual driving, the system achieved 0.27 m mean total error and 0.20 m mean cross-track error. However, in autonomous mode, the mean total error went up to 0.5 m and the mean cross-track error was 0.36 m. (Note that this includes all weather conditions.) The degraded performance in autonomous mode can be attributed to the fact that the autonomous systems do not follow the mapped path perfectly. However, the magnitude of the error was not so significant to preclude following the reference path and over the course of over 8 km of autonomous testing not a single unplanned disengagement was necessary.

VII. Conclusion

This letter presents a solution for autonomous driving using only underground features measured with an LGPR sensor.

This solution permits safely driving in inclement weather including rain, and snow, and under complete darkness at night. Furthermore, we have obtained measurements of the system performance during fully autonomous navigation, which, to the best of our knowledge, has never been published previously. While safety requirements at our testing facility did not allow for testing autonomously at high speeds, [3] tested a similar setup (albeit with a human driver) at highways speeds and demonstrated successful localization. Therefore, we believe our results will readily extend to high speed driving as well. Current limitations of the system involve a decrease in performance when the weather conditions during localization differ from those during map creation. Furthermore, the current system is not capable of performing global localization without using a GPS prior. Thus, future work will investigate methods explicitly modeling temperature and humidity induced reflection changes as well as aggregation-based global localization techniques for precise initialization even in the absence of GPS priors.

REFERENCES

[1] C. Cadena *et al.*, "Past, present, and future of simultaneous localization and mapping: Toward the robust-perception age," *IEEE Trans. Robot.*, vol. 32, no. 6, pp. 1309–1332, Dec. 2016.

[2] S. Lowry *et al.*, "Visual place recognition: A survey," *IEEE Trans. Robot.*, vol. 32, no. 1, pp. 1–19, Feb. 2016.

[3] M. Cornick, J. Koechling, B. Stanley, and B. Zhang, "Localizing ground penetrating RADAR: A step toward robust autonomous ground vehicle localization," *J. Field Robot.*, vol. 33, no. 1, pp. 82–102, 2016.

[4] J. Dickmann *et al.*, "Automotive radar the key technology for autonomous driving: From detection and ranging to environmental understanding," in *Proc. IEEE Radar Conf.*, 2016, pp. 1–6.

[5] R. H. Rasshofer and K. Gresser, "Automotive radar and lidar systems for next generation driver assistance functions," in *Proc. Adv. Radio Sci.*, vol. 3, 2005, pp. 205–209.

[6] K. Mori, T. Takahashi, I. Ide, H. Murase, T. Miyahara, and Y. Tamatsu, "Recognition of foggy conditions by in-vehicle camera and millimeter wave radar," in *Proc. IEEE Intell. Veh. Symp.*, 2007, pp. 87–92.

[7] D. T. Linzmeier, M. Skutek, M. Mekhaiel, and K. C. J. Dietmayer, "A pedestrian detection system based on thermopile and radar sensor data fusion," in *Proc. 7th Int. Conf. Inf. Fusion*, 2005, pp. 1272–1279.

[8] M. E. Bouzouraa and U. Hofmann, "Fusion of occupancy grid mapping and model based object tracking for driver assistance systems using laser and radar sensors," in *Proc. IEEE Intell. Veh. Symp.*, 2010, pp. 294–300.

[9] D. Vivet, F. Gérossier, P. Checchin, L. Trassoudaine, and R. Chapuis, "Mobile ground-based radar sensor for localization and mapping: An evaluation of two approaches," *Int. J. Adv. Robot. Syst.*, vol. 10, no. 8, p. 307, 2013.

[10] F. Schuster, M. Wörner, C. G. Keller, M. Haueis, and C. Curio, "Robust localization based on radar signal clustering," in *Proc. IEEE Intell. Veh. Symp.*, 2016, pp. 839–844.

[11] E. Ward and J. Folkesson, "Vehicle localization with low cost radar sensors," in *Proc. IEEE Intell. Veh. Symp.*, 2016, pp. 864–870.

[12] B. Bescos, J. M. Facil, J. Civera, and J. Neira, "DynaSLAM: Tracking, Mapping, and Inpainting in Dynamic Scenes," *IEEE Robot. Autom. Lett.*, vol. 3, no. 4, pp. 4076–4083, Oct. 2018.

[13] B. Bescos, J. Neira, R. Siegwart, and C. Cadena, "Empty cities: Image inpainting for a dynamic-object-invariant space," in *Proc. Int. Conf. Robot. Autom.*, 2019, pp. 5460–5466.

[14] M. Dymczyk, E. Stumm, J. Nieto, R. Siegwart, and I. Gilitschenski, "Will it last? Learning stable features for long-term visual localization," in *Proc. 4th Int. Conf. 3D Vision*, 2016, pp. 572–581.

[15] M. Bürki, C. Cadena, I. Gilitschenski, R. Siegwart, and J. Nieto, "Appearance-based landmark selection for visual localization," *J. Field Robot.*, vol. 36, no. 6, pp. 1041–1073, 2019.

[16] W. Churchill and P. Newman, "Experience-based navigation for long-term localisation," *Int. J. Robot. Res.*, vol. 32, no. 14, pp. 1645–1661, 2013.

[17] J. L. Davis and A. Annan, "Ground-penetrating radar for high-resolution mapping of soil and rock stratigraphy," *Geophysical Prospecting*, vol. 37, pp. 531–551, 1989.

[18] J. Rea and R. Knight, "Geostatistical analysis of ground-penetrating radar data: A means of describing spatial variation in the subsurface," *Water Resour. Res.*, vol. 34, no. 3, pp. 329–339, 1998.

[19] S.-H. Ni, Y.-H. Huang, K.-F. Lo, and D.-C. Lin, "Buried pipe detection by ground penetrating radar using the discrete wavelet transform," *Comput. Geotechnics*, vol. 37, no. 4, pp. 440–448, 2010.

[20] J. J. Daniels *et al.*, "Ground penetrating radar for imaging archeological objects," in *Proc. New Millennium Int. Forum Conservation Cultural Property*, 2000, pp. 247–265.

[21] K. M. Dawson-Howe and T. G. Williams, "The detection of buried landmines using probing robots," *Robot. Auton. Syst.*, vol. 23, pp. 235–243, 1998.

[22] R. J. Chignell, "Ground penetrating radar-a sensor for mine detection," in *Proc. EUREL Int. Conf. Detection Abandoned Landmines*, 1996, pp. 103–108.

[23] R. M. Williams, L. E. Ray, and J. Lever, "An autonomous robotic platform for ground penetrating radar surveys," in *Proc. IEEE Int. Geosci. Remote Sens. Symp.*, 2012, pp. 3174–3177.

[24] J. H. Lever, A. J. Delaney, L. E. Ray, E. Trautmann, L. A. Barna, and A. M. Burzynski, "Autonomous GPR surveys using the polar rover yeti," *J. Field Robot.*, vol. 30, no. 2, pp. 194–215, 2013.

[25] H. M. La *et al.*, "Autonomous robotic system for high-efficiency non-destructive bridge deck inspection and evaluation," in *Proc. IEEE Int. Conf. Autom. Sci. Eng.*, 2013, pp. 1053–1058.

[26] Y. Su *et al.*, "The preliminary results of lunar penetrating radar on board the chinese chang'e-3 rover," in *Proc. 15th Int. Conf. Ground Penetrating Radar*, 2014, pp. 1032–1036.

[27] E. Jose and M. D. Adams, "An augmented state SLAM formulation for multiple line-of-sight features with millimetre wave RADAR," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2005, pp. 3087–3092.

[28] T. Ort *et al.*, "Maplite: Autonomous intersection navigation without a detailed prior map," *IEEE Robot. Autom. Lett.*, vol. 5, no. 2, pp. 556–563, Apr. 2019.

[29] "Federal highway administration," 2019. [Online]. Available: https://safety.fhwa.dot.gov/geometric/pubs/mitigationstrategies/chapter3/3_lanewidth.cfm

[30] S. Kennedy and J. Rossi, "Performance of a deeply coupled commercial grade GPS/INS system from KVH and NovAtel Inc," in *Proc. Position, Location Navigation Symp.*, 2008, pp. 17–24.

[31] E. W. Dijkstra, "A note on two problems in connexion with graphs," *Numerische Mathematik*, vol. 1, pp. 269–271, 1959.

[32] T. Saarenketo and T. Scullion, "Road evaluation with ground penetrating radar," *J. Appl. Geophys.*, vol. 43, pp. 119–138, 2000.

[33] R. E. Kalman, "A new approach to linear filtering and prediction problems," *J. Basic Eng.*, vol. 82, pp. 35–45, 1960.

[34] S.-L. Sun and Z.-L. Deng, "Multi-sensor optimal information fusion kalman filter," *Automatica*, vol. 40, no. 6, pp. 1017–1023, 2004.

[35] G. Bishop *et al.*, "An introduction to the Kalman filter," in *Proc. ACM Special Interest Group Comput. Graph.*, vol. 8, no. 27599-23175, p. 41, 2001.

[36] T. Moore and D. Stouch, "A generalized extended Kalman filter implementation for the robot operating system," *Intell. Auton. Syst.*, vol. 13, pp. 335–348, 2014.

[37] "Oxts rt user manual," 2019. [Online]. Available: https://www.oxts.com/app/uploads/2018/02/rtman.pdf

[38] C. F. F. Karney, "Algorithms for geodesics," *J. Geodesy*, vol. 87, no. 1, pp. 43–55, 2013.

[39] R. C. Coulter, "Implementation of the pure pursuit path tracking algorithm," Carnegie-Mellon Univ., Pittsburgh, PA, USA, Tech. Rep. CMU-RI-TR-92-01, 1992.

[40] J. Leonard *et al.*, "Team mit urban challenge technical report," MIT Comput. Sci. Artif. Intell. Lab., Cambridge, MA, USA, Tech. Rep. MIT-CSAIL-TR-2007-058, 2007.