

# Ground Encoding: Learned Factor Graph-based Models for Localizing Ground Penetrating Radar

Alexander Baikovitz<sup>1</sup>, Paloma Sodhi<sup>1</sup>, Michael Dille<sup>2</sup>, Michael Kaess<sup>1</sup>

**Abstract**— We address the problem of robot localization using ground penetrating radar (GPR) sensors. Current approaches for localization with GPR sensors require *a priori* maps of the system’s environment as well as access to approximate global positioning (GPS) during operation. In this paper, we propose a novel, real-time GPR-based localization system for unknown and GPS-denied environments. We model the localization problem as an inference over a factor graph. Our approach combines 1D single-channel GPR measurements to form 2D image submaps. To use these GPR images in the graph, we need sensor models that can map noisy, high-dimensional image measurements into the state space. These are challenging to obtain *a priori* since image generation has a complex dependency on subsurface composition and radar physics, which itself varies with sensors and variations in subsurface electromagnetic properties. Our key idea is to instead learn relative sensor models directly from GPR data that map non-sequential GPR image pairs to relative robot motion. These models are incorporated as factors within the factor graph with relative motion predictions correcting for accumulated drift in the position estimates. We demonstrate our approach over datasets collected across multiple locations using a custom designed experimental rig. We show reliable, real-time localization using only GPR and odometry measurements for varying trajectories in three distinct GPS-denied environments.

## I. INTRODUCTION

We focus on the problem of localization using ground penetrating radar (GPR) in unknown and GPS-denied environments. Hardware failure, repetitive or sparse features, and poor visibility and illumination can make localization in warehouses, mines, caves and other enclosed, unstructured environments challenging. Consider an operational subsurface mine where continuous drilling and blasting changes the line-of-sight appearance of the scene and creates unexplored environments, which could lead to poor localization using visual and spatial information. While the visual environment may change, subsurface features are typically invariant and can be used to recognize the system’s location.

Currently, GPR is widely used in utility locating, concrete inspection, archaeology, unexploded ordinance and landmine identification, among a growing list of applications to determine the depth and position of subsurface assets [1, 2]. GPR information is often difficult to interpret because of noise, variable subsurface electromagnetic properties, and sensor variability over time [3]. A single GPR measurement only reveals limited information about a system’s environment,

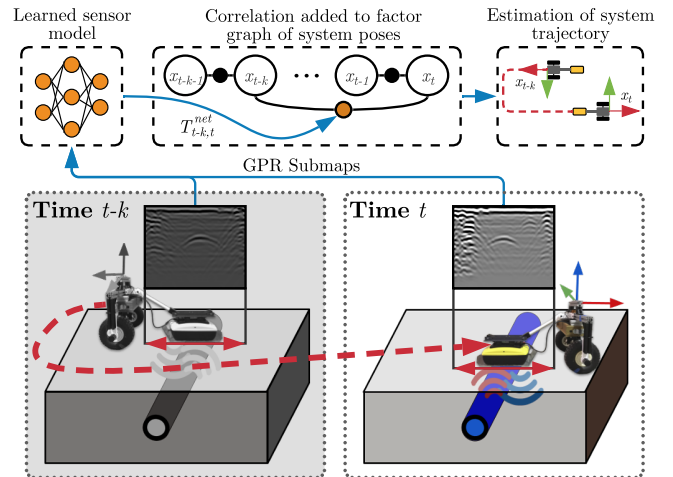


Fig. 1. Estimating poses for a ground vehicle using subsurface measurements from Ground Penetrating Radar (GPR) as inference over a factor graph. Since GPR measurements are challenging to correlate, a relative transformation is learned between submaps to correct the system’s pose.

requiring a sequence of measurements to discern the local structure of the subsurface. In order to use GPR for localization, we need a representation for GPR data that (1) captures prominent local features and (2) is invariant over time.

In GPS-denied environments, operating using only proprioceptive information (e.g. IMU, wheel encoders) will accumulate drift. Prior work [4, 5] on GPR localization has addressed this challenge by using an *a priori* global map of arrayed GPR images. During operation, the system compares the current GPR image against this global map database, and corrects for accumulated drift within a filtering framework. However, this approach requires approximate global positioning and a prior map of the operating environment which would not generalize to previously unknown environments.

We propose an approach that allows for correction of this accumulated drift without any prior map information. In the absence of a prior map, we must reason over multiple GPR measurements together to be able to infer the latent robot location. We formulate this inference problem using a factor graph, which is now common with many modern localization and SLAM objectives [6–10]. GPR, inertial, and wheel encoder measurements are incorporated into the graph as factors to estimate the system’s latent state, consisting of position, orientation, velocity, and IMU biases. To incorporate measurements into the graph, a sensor model is needed to map measurements to states. For GPR sensors,

This work was supported by a NASA Space Technology Graduate Research Opportunity.

<sup>1</sup> The Robotics Institute, Carnegie Mellon University, <sup>2</sup> Intelligent Robotics Group, NASA Ames Research Center, Correspondence to abaikovitz@cmu.edu

*a priori* models are typically challenging to obtain since image generation has a complex dependency on subsurface composition and radar physics. Instead, we learn relative sensor models that map non-sequential GPR image pairs to relative robot motion. The relative motion information in turn enables us to correct for drift accumulated when using just proprioceptive sensor information. Our main contributions are:

- 1) A formulation of the GPR localization problem as inference over a factor graph without a prior map.
- 2) A learnable GPR sensor model based on submaps.
- 3) Experimental evaluation on a test platform with a single-channel GPR system operating in three different GPS-denied environments.

## II. RELATED WORK

**GPR in robotics:** Most use of GPR in the robotics domain has been for passive inspection, which includes excavation of subsurface objects [11], planetary exploration [12–14], mine detection [15, 16], bridge deck inspection [17], and crevasse identification [18].

Localizing GPR (LGPR) was introduced by Cornick *et al.* to position a custom GPR array in a prior map [4]. In this method, the current GPR measurement is registered to a grid of neighboring prior map measurements using a particle swarm optimization, which attempts to find the 5-DOF position of the measurement that maximizes the correlation with the grid data. Ort *et al.* extended this work using the same LGPR device for fully autonomous navigation without visual features in variable weather conditions [5]. This approach involves two Extended Kalman Filters; one filter estimates the system velocities from wheel encoder and inertial data and the other filter fuses these velocities with the LGPR-GPS correction. In the evaluation across different weather conditions, it is apparent that GPR measurements acquired in rain and snow were less correlated to the *a priori* map than in clear weather using the heuristic correlation for registration method proposed in [4].

**Factor graphs for SLAM:** Modern visual and spatial localization systems often use smoothing methods, which have been proven to be more accurate and efficient than classical filtering-based methods [7]. The smoothing objective is typically framed as a MAP inference over a factor graph, where the variable nodes represent latent states and factor nodes represent measurement likelihoods. To incorporate sensor measurements into the graph, a sensor model is needed to map high dimensional measurements to a low-dimensional state. Typically these sensor models are analytic functions, such as camera geometry [19, 20] and scan matching [21].

**Learned sensor models:** Learning-based methods provide an alternate option to model complex sensor measurements. Prior work in visual SLAM has produced dense depth reconstructions from learned feature representations of monocular camera images [8, 10]. In computer vision, spatial correlation networks have been used to learn optical flow and localize RGB cameras in depth maps [22–24]. Recent work on object state estimation using tactile feedback has demonstrated

the effectiveness of learned sensor models in factor graph inference [9].

Our approach enables robust GPR-based positioning by leveraging the benefits of structured prediction in factor graph inference with learning-based sensor models. We use a smoothing-based approach, which is more accurate and efficient than filtering-based approaches used by previous localizing GPR systems, and can be solved in real-time using state-of-the-art algorithms [25, 26]. Learning-based sensor models for GPR can outperform engineered correlation approaches by exploiting unique features in a sequence of GPR measurements, enabling improved performance in even sparsely featured environments.

## III. PROBLEM FORMULATION

We formulate our GPR localization problem as inference over a factor graph. A factor graph is a bipartite graph with two types of nodes: variables  $x \in \mathcal{X}$  and factors  $\phi(\cdot) : \mathcal{X} \rightarrow \mathbb{R}$ . Variable nodes are the latent states to be estimated, and factor nodes encode constraints on these variables such as measurement likelihood functions.

Maximum a posteriori (MAP) inference over a factor graph involves maximizing the product of all factor graph potentials, i.e.,

$$\hat{x} = \operatorname{argmax}_x \prod_{t=1}^T \phi_t(x) \quad (1)$$

Under Gaussian noise model assumptions, MAP inference is equivalent to solving a nonlinear least-squares problem [6]. That is, for Gaussian factors  $\phi_t(x)$  corrupted by zero-mean, normally distributed noise,

$$\begin{aligned} \phi_t(x) &\propto \exp \left\{ -\frac{1}{2} \|f_t(x) - z_t\|_{\Sigma_t}^2 \right\} \\ \Rightarrow \hat{x} &= \operatorname{argmin}_x \sum_{t=1}^T \|f_t(x) - z_t\|_{\Sigma_t}^2 \end{aligned} \quad (2)$$

where,  $f_t(x)$  is the measurement likelihood function predicting expected measurement given current state,  $z_t$  is the actual measurement, and  $\|\cdot\|_{\Sigma_t}$  is the Mahalanobis distance with measurement covariance  $\Sigma_t$ .

For the GPR localization problem, variables in the graph at time step  $t = 1 \dots T$  are the 6-DOF robot poses  $s_t \in SE(3)$ , velocities  $v_t$ , and IMU biases  $b_t$ , i.e.  $x_t = [s_t \ v_t \ b_t]^T$ . Factors in the graph incorporate different likelihoods for GPR, IMU, and wheel encoder measurements. At every time step  $t$ , new variables and factors are added to the graph. Writing out Eq. 2 for the GPR localization objective,

$$\begin{aligned} \hat{x}_{1:T} = \operatorname{argmin}_{x_{1:T}} \sum_{t=1}^T \left\{ \|f_{gpr}(x_{t-k}, x_t) - z_{t-k,t}^{gpr}\|_{\Sigma_{gpr}}^2 + \right. \\ \left. \|f_{wh}(x_{t-1}, x_t) - z_{t-1,t}^{wh}\|_{\Sigma_{wh}}^2 + \right. \\ \left. \|f_{imu}(x_{t-1}, x_t) - z_{t-1,t}^{imu}\|_{\Sigma_{imu}}^2 \right\} \end{aligned} \quad (3)$$

This objective in Eq. 3 is solved online every time step using an efficient, incremental solver iSAM2 [26]. Individual terms in Eq. 3 are described in more detail in Section IV.

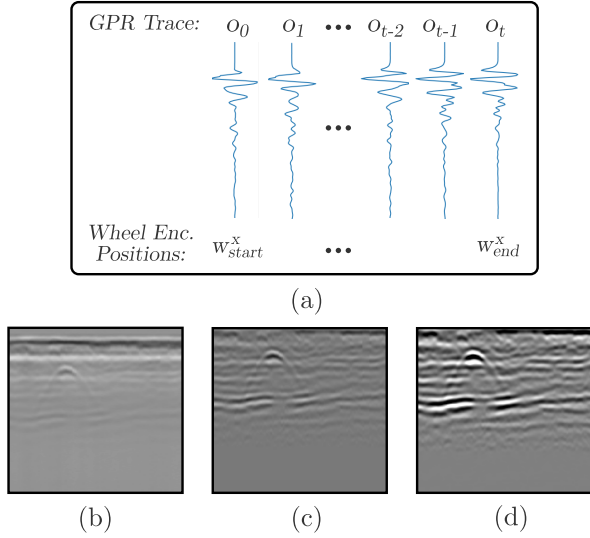


Fig. 2. Overview of GPR image construction. (a) Unprocessed localized traces received by the device. (b) Horizontally stacked traces from (a), where the amplitudes correspond to pixel intensity. (c) Measurements after filtering and gain. (d) Final image after thresholding.

#### IV. APPROACH

There are three primary stages in our approach: GPR preprocessing and submapping (Section IV-A), learning GPR models (Section IV-B), and factor graph optimization (Section IV-C). Fig. 5 illustrates our ground encoding approach.

##### A. GPR Submapping

**Preprocessing:** Digital signal processing techniques are needed to improve the signal-to-noise ratio of the received GPR signal, counteract signal attenuation, and remove recurrent artifacts that are intrinsic to GPR systems. Measurements  $o_t$  arrive from each position as a 1D discrete waveform called a trace, which are locally averaged to reduce noise and resampled to form a uniformly spaced image [27].

Receiving signals often contain a low frequency component and DC bias caused by saturation and inductive coupling, requiring a dewow filter, which involves a DC subtraction and low-cut filter parameterized on the radar's center frequency and bandwidth [1]. In order to counteract attenuation, we multiply the signal by a Spreading and Exponential Compensation (SEC) gain function,

$$G = \exp(a \cdot t) \cdot t^b \quad (4)$$

where,  $a$  is the exponential gain constant and  $b$  is the power gain constant [28].

GPR images have global and local context. Global features such as horizontal layers encode information about the general environment of a system, which can include boundaries like layers of asphalt in roads. Local features, which include pipes and point objects, provide salient features needed to effectively localize. To maintain prominent global features, we subtract the mean trace over all prior submaps (as

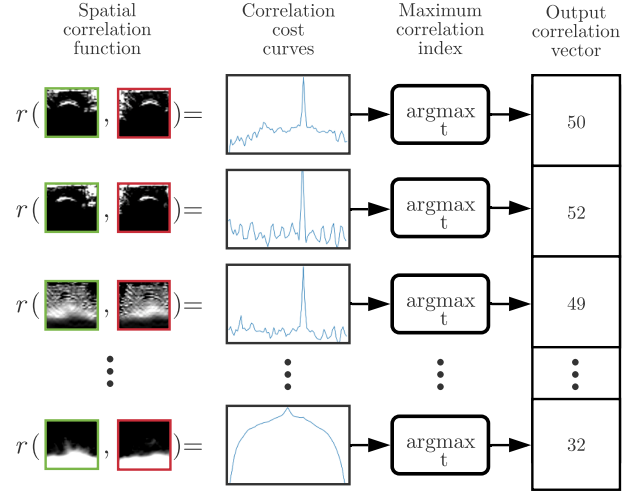


Fig. 3. Summary of spatial correlation layer for submap transform prediction. Feature activation maps are provided to a spatial correlation function, which computes a curve of image correlations with translation. The maximum values of each curve are compiled into a correlation vector.

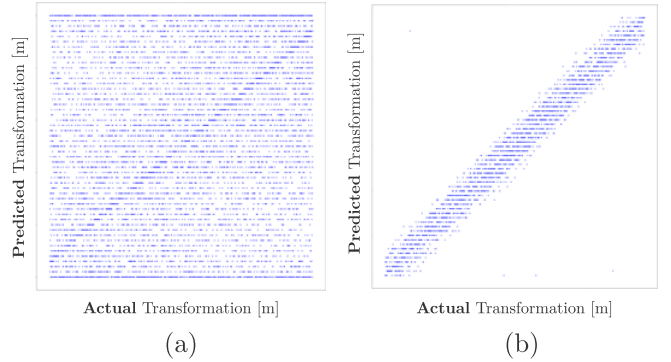


Fig. 4. Actual versus predicted transformation (a) without spatial correlation and (b) with spatial correlation. The spatial correlation layer converts it into a linear relationship, simplifying the learning process for the transform prediction network.

opposed to over each submap) to remove repeated artifacts introduced by the radar while emphasizing local features.

**Submap formation:** Traditional GPR systems provide a one-dimensional trace  $o_t$  at each position, which does not provide enough information to effectively determine a system's unique position. A collection of traces represents a local fingerprint that may contain valuable information for localization. Our approach involves constructing submaps  $S_t$  based on integration of local wheel encoder measurements  $w_t^x$  (Fig. 2). A submap is approximated by sampling discrete GPR measurements from a continuous distribution of GPR measurements using an interpolation function  $\hat{F}$  to create a uniformly sampled image,

$$S_t = \hat{F}(\{\bar{o}_t, w_t^x \mid w_{start}^x \leq w_t^x < w_{end}^x\}) \quad (5)$$

where,  $\bar{o}_t$  is the locally averaged GPR measurement, de-

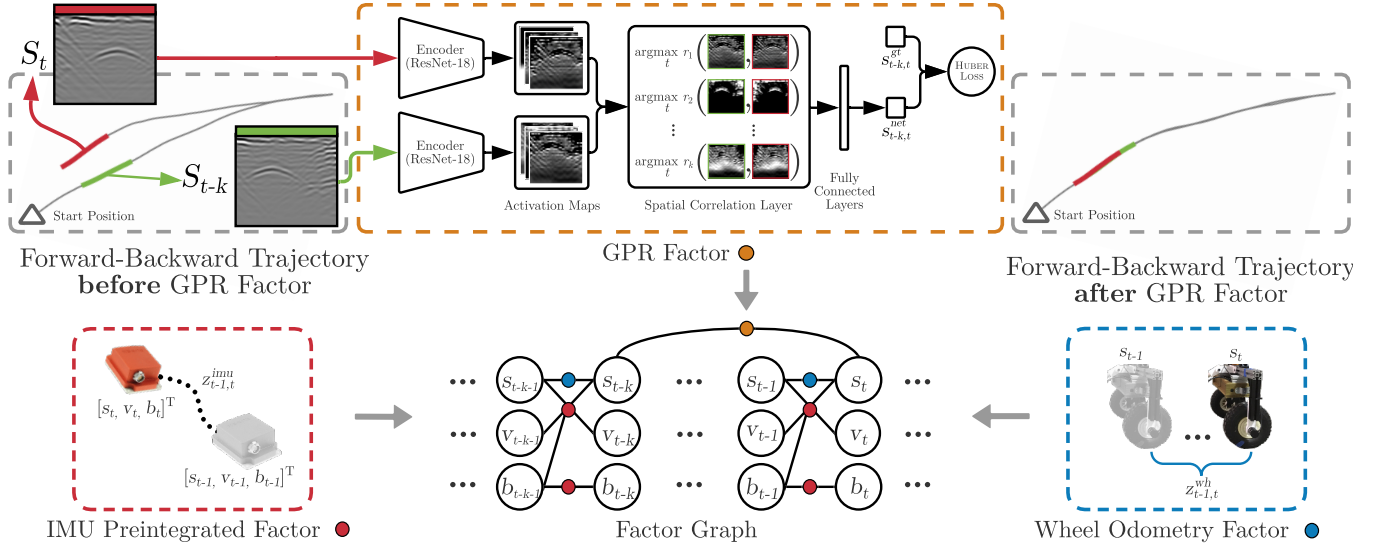


Fig. 5. Overview of the Ground Encoding system. On the left, two submaps containing similar subsurface structures are acquired during forward-backward trajectory that has accumulated drift. An autoencoder maps these submaps to feature activation maps, which are used by the transform prediction network to find the relative transformation. Factors denoting the measurement likelihoods from GPR, inertial, and wheel encoder measurements are incorporated into the graph to estimate the system’s latent states  $x_t = [s_t, v_t, b_t]^T \forall t \in \{1 \dots T\}$ . The trajectory on the right shows the outcome of the GPR correction.

scribed in Preprocessing, at time  $t$  and  $w_t^x$  is the wheel encoder measurement at time  $t$ . We use a rule-based method that maps a set of wheel odometry and angular velocity measurements to a decision of whether a short trajectory is a valid submap.

### B. Learned GPR model

The goal of the learned GPR sensor model is to derive the GPR factor cost term in Eq. 3 for the factor graph optimization. The sensor model is trained to predict a relative transformation between two submaps  $\{S_{t-k}, S_t\}$  from ground truth measurements. Submap images are compressed into feature activation map, which are provided to a transform prediction network to estimate the relative transformation  $s_{t-k,t}^{net}$  between poses  $s_{t-k}$  and  $s_t$  (Fig. 5).

**Feature learning:** Identifying the function that directly maps submaps to transformations is prone to overfitting because of substantial noise in the original radar submaps. To learn the function that relates two similar submaps  $S_{t-k}$  and  $S_t$ , an intermediate feature representation  $f_{t-k}$  and  $f_t$  is needed. A ResNet-18 auto-encoder is used to obtain  $k$  feature activation maps that contain relevant features for localization like vertical edges. The autoencoder is trained with an  $L1$  reconstruction loss to preserve the sparsity of features in the original submap.

**Submap comparison:** Prior to identifying transformations in the data, we perform an average pooling on the row-wise standard deviation of submap images to check for salient features. If submaps  $S_{t-k}$  and  $S_t$  contain valid features, a linear correlation network checks whether they share common features using a method further described in Spatial Correlation and Transform Prediction.

**Spatial Correlation:** If submaps contain salient features and are correlated, we predict a relative 1D transformation  $s_{t-k,t}^{net}$  using a two stage approach. We first construct a set of cost curves by comparing each feature map  $S_{t-k,i}$  and  $S_{t,i}$  for all  $k$  using Eq. 9 for feature activation map pixel space. We then evaluate the argmax of each cost curve, which is encoded in a vector as shown in Fig. 3.

Modeling the transformation directly between two images tended to overfit to training data and produced incorrect transform predictions as shown in Fig. 4(a). The spatial correlation layer simplifies transform prediction by naturally embedding an argmax in the learning process, reducing the learning problem to linear regression as shown in Fig. 4(b). This approach is generally relevant for predicting transformations between images and volumes for GPR and likely other sensing modalities.

**Transform prediction:** The transform prediction network learns weights for different filters and a scale factor to convert pixel space to metric range space (e.g. distance in meters). The spatial correlation vector is provided to a fully-connected regression network to identify the relative transformation  $s_{t-k,t}^{net}$ . Then, the network learns to provide higher weight to filters that are better correlated with the desired translation, reducing the effect of filters that encode common patterns in the data like horizontal banding. The transform prediction network is trained using Huber loss against supervised ground truth data  $s_{t-k,t}^{gt}$  acquired by a robotic total station.

### C. Factor graph optimization

**GPR factor:** Measurements from GPR are incorporated into the graph as factors with the cost from Eq. 3. The relative GPR factor has a quadratic cost penalizing large residual



values and is defined as,

$$\|f_{gpr}(x_{t-k}, x_t) - z_{t-k,t}\|_{\Sigma_{gpr}}^2 := \|s_t^{-1} s_{t-k} \ominus z_{t-k,t}^{gpr}\|_{\Sigma_{gpr}}^2 \quad (6)$$

where,  $z_{t-k,t}^{gpr}$  is the predicted relative transformation from the transform prediction network,  $s_t^{-1} s_{t-k}$  is the estimated relative pose transformation between two variable nodes in the graph, and  $\ominus$  represents the difference between two manifold elements.

**IMU preintegrated factor:** For incorporating IMU measurements in the graph, we use the IMU preintegration factor proposed in [29]. The preintegration factor locally integrates multiple high-frequency accelerometer and gyroscope measurements into a single compound preintegrated IMU measurement. This has the advantage of combining the speed and complexity benefits of filtering-based methods along with the accuracy of smoothing methods.

The IMU factor term from Eq. 3 can be expressed as a residual over the differences in orientation  $\Delta R_{i,j}$ , velocity  $\Delta v_{i,j}$ , and position  $\Delta t_{i,j}$ ,

$$\|f_{imu}(x_{t-1}, x_t) - z_{t-1,t}\|_{\Sigma_{imu}}^2 = \|\mathbf{r}_{\mathcal{I}_{ij}}\|^2 + \|\mathbf{r}_{b_{ij}}\|^2 \quad (7)$$

where,  $\{i, j\}$  are state index pairs between which preintegration is performed. We perform the preintegration between consecutive states, i.e.  $\{i, j\} := \{t-1, t\}$ . Here,  $\mathbf{r}_{\mathcal{I}_{ij}} = [\mathbf{r}_{\Delta R_{i,j}}^T, \mathbf{r}_{\Delta v_{i,j}}^T, \mathbf{r}_{\Delta t_{i,j}}^T]^T$  is the preintegration error residual and  $\mathbf{r}_{b_{ij}}$  is bias term estimation errors. We refer the reader to [29] for more details on these residual error terms.

**Wheel encoder factor:** Wheel encoder measurements are incorporated into the graph between sequential poses  $\{s_{t-1}, s_t\}$ . The relative wheel encoder factor is defined as,

$$\|f_{wh}(x_{t-1}, x_t) - z_{t-1,t}^{wh}\|_{\Sigma_{wh}}^2 := \|s_t^{-1} s_{t-1} \ominus z_{t-1,t}^{wh}\|_{\Sigma_{wh}}^2 \quad (8)$$

where,  $z_{t-k,t}^{wh}$  is the relative difference between two poses measured by the wheel encoder.

## V. RESULTS AND EVALUATION

We evaluate our GPR-based localization system in three distinct, GPS-denied environments: a basement (*nsh\_b*), a factory floor (*nsh\_h*), and a parking garage (*gates\_g*). We first discuss our experimental setup in Section V-A. We then provide an ablation of validation losses from different learned and engineered GPR sensor models in Section V-C. We finally evaluate our entire GPR-based localization system and demonstrate the effects of adding GPR information in the graph optimization in Section V-D. Data and utility functions used for this evaluation are provided in the CMU-GPR dataset<sup>1</sup> [30].

### A. Experimental setup

We constructed a manually-pulled test rig named SuperVision shown in Fig. 6(a) for data acquisition. SuperVision uses a quad-core Intel NUC for compute and wirelessly transmits onboard data from an XSENS MTI-30 9-axis

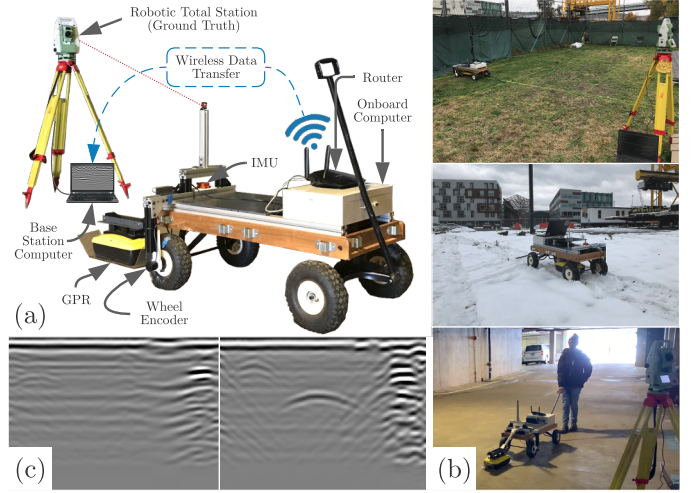


Fig. 6. (a) Experimental setup for ground encoding datasets. (b) Testing in distinct environments. (c) Left GPR image constructed with 2.5cm antenna-ground separation. Right GPR image constructed with ground-coupled antenna configuration.

Inertial Measurement Unit, YUMO quadrature encoder with 1024 PPR, and a Sensors and Software Noggin 500 GPR. Readings from the IMU magnetometer were excluded due to intermittent magnetic interference commonly found in indoor environments. Ground truth data was acquired by a Leica TS15 total station. The base station logs measurements from the onboard computer and the total station to ensure consistent timing.

In our testing, we found that GPR antenna placement was a critical design decision. Our initial design suspended the radar system 2.5cm above the ground to maintain a fixed transformation between the IMU and GPR sensor. While a fixed sensor setting like this is most common for localization systems, the small air gap between the GPR and the ground introduced ground energy losses and multi-path interactions causing poor depth penetration and repetitive ringing. To address this, we added a passive suspension to the GPR system so as to maintain constant ground contact. This improved depth penetration as shown in the right image of Fig. 6(c). We address methods to improve the robustness and applicability of our system in Section VI.

### B. Engineered GPR model baseline

For evaluating our approach, we also consider a baseline *engineered* GPR model that makes use of a correlation metric similar to existing work. This *engineered* GPR model is incorporated as factors in the graph in the same way as our learned GPR model.

Existing GPR-based localization methods use a measurement correlation value to register images to a prior map [4, 5]. We compute a similar measurement correlation value between image pairs and select the pair with the highest correlation to add as factors in the graph. We use the Pearson product-moment correlation between two non-

<sup>1</sup>The CMU-GPR dataset is available at <https://github.com/rpl-cmu/CMU-GPR-Dataset>




Fig. 7. Qualitative evaluation of vehicle trajectory estimate aligned with ground truth trajectory. (a) The robot follows a trajectory that loops over itself twice. (b) The robot starts moving on the vertical segment, then moves along the horizontal segment, before finally revisiting the vertical segment. In both cases, the learned sensor model follows the ground truth data better than the engineered method despite receiving the same loop closure detections (qualitative oracle results not shown for improved clarity).




Fig. 8. RMSE Absolute Trajectory Error (ATE) for each testing environment.




Fig. 9. RMSE ATE over time. (a) Corresponds to the trajectory from Fig. 7(a). (b) Corresponds to the trajectory from Fig. 7(b). In both cases, loop closure events reduce the accumulation of odometry error.

sequential GPR images  $S_{t-k}, S_t$ ,

$$r(A(T), B(T)) = \frac{\sum_{r,c} A_{r,c}(T) B_{r,c}(T)}{\sqrt{\sum_{r,c} A_{r,c}(T)^2} \sqrt{\sum_{r,c} B_{r,c}(T)^2}} \quad (9)$$

where,  $A_{r,c}(T) = S_{t-k}(T, r, c) - \text{mean}(S_{t-k}(T))$ ,  $B_{r,c}(T) = S_t(T, r, c) - \text{mean}(S_t(T))$ ,  $r$  is the submap row,

$c$  is the submap column, and  $T$  is the transformation from  $S_{t-k}$  to  $S_t$  used to only compare the shared region between the submaps.

Specifically, we estimate the transformation along the x-direction of the robot's motion. This is obtained by solving the optimization that maximizes the correlation from Eq. 9,

$$z_{t-k,t} = K \arg\max_T (r(A(T), B(T))) \quad (10)$$

where,  $K$  is a constant that converts submap pixel space to robot motion space.

### C. GPR factor learning

We now evaluate the performance of different GPR models to be used in the factor graph. Table I compares losses for the baseline *engineered* models against our learned GPR sensor model discussed in Section IV-B for different choices of network architectures. The loss here is the mean-squared error that we saw in Eq. 6 against ground truth transformations. We use the Huber loss function for robustness to outliers. *engineered* is the baseline correlation approach discussed in Section V-C which resembles prior work [4]. *zeroth* is a zeroth-order model that predicts the average relative

TABLE I  
VALIDATION LOSSES FOR DIFFERENT GPR MODELS (CM)

Model type	Dataset			Combined
	gates_g	nsh_b	nsh_h	
engineered	97.66	88.07	100.24	95.33
zeroth	17.81	137.97	23.30	59.49
<b>linear, corr-feat</b>	<b>14.81</b>	<b>2.88</b>	<b>3.90</b>	<b>7.40</b>
nonlinear, corr-feat	14.96	3.09	4.06	8.97

transform of the training dataset. *corr-feat* are acquired by concatenating the argmax values of the discrete cost curves from the spatial correlation of autoencoder activation maps as described in Section IV-B. *linear* and *nonlinear* refer to the linear and nonlinear activation functions in the fully-connected layers.

The *engineered* model has notably high losses as there can be many false positive matches since the GPR image data often has similar amplitude with subtle features. In comparison, we see that the different learned GPR model architectures have much lower losses, with *linear, corr-feat* model type having the best performance. We also found that spatial correlation features *corr-feat* generalized much better than using vectorized autoencoder feature maps directly.

#### D. Factor graph optimization

We finally demonstrate the effect of the GPR sensor model on odometry drift correction through a qualitative and quantitative evaluation. The best performing model from Table I, *linear, corr-feat*, was selected as the *learned* model in this evaluation. We compare the *engineered* and *learned* model with the ground truth trajectory and the results from an *oracle* model. The *oracle* model predicts the ground truth relative pose transformation, representing the lower error bound of the GPR factor. The covariance parameters in the graph are fixed and the same across sequences.

**Qualitative Evaluation:** Both trajectories in Fig. 7 were collected in a modern parking garage with homogeneous and repetitious subsurface features, making the test challenging for our GPR system. Additionally, non-deterministic reflections from nearby cars and walls were present in the dataset. As shown, the sensor model relying on *engineered* features identifies incorrect transformations by comparing processed submaps directly, causing substantial drift. This is also reflected in Fig. 9, where incorrect loop closures cause the *engineered* model to have greater error than odometry alone. We see that the *learned* model recovers the robot’s poses close to the true trajectory and matches the *oracle* performance closely.

**Quantitative Evaluation:** Fig. 8 shows the RMSE of the absolute trajectory error described in [31]. Errors were computed over 7 sequences in *gates\_g*, 5 sequences in *nsh\_b*, and 3 sequences in *nsh\_h* (15 sequences overall). Odometry measurements had greater error in the *gates\_g* set since loops and turns were more common. To better compare sensor models, an oracle provided identical loop closure observations to both *engineered* and *learned* models.




Fig. 10. Optimizer runtimes for entire set of *nsh\_b* sequences.

The *learned* model outperformed the *engineered* model in all sequences. The *engineered* model would often produce false detections in sparsely featured environments, causing a large variance in performance. Repetitive structures and noise in the processed image caused the *engineered* model to perform poorly. The *learned* sensor model performed consistently better by decomposing the image into feature activation maps, containing prominent structures and edges in the data that are easier to compare. The *learned* model nearly performs as well as the *oracle* detector, the model’s lower error bound.

**Runtime performance:** The optimizer runtime remains relatively constant with the addition of new measurements and GPR constraints to the graph. Fig. 10 shows real-time performance for estimating 2000 variables over 22 minutes of operation.

## VI. DISCUSSION

We presented a GPR-based localization system that uses a learned sensor model to position a robot in a unknown, GPS-denied environment. We accomplished this by combining GPR measurements to create individual submaps, which were provided to a transform prediction network to predict a relative pose transformation. These transformations were then incorporated as measurement likelihoods in a factor graph, where the GPR measurement corrected for accumulated drift from proprioceptive measurements. We validated our system in three distinct environments, where we have shown improved localization performance over existing correlation-based approaches.

Prior work has relied on sophisticated multi-channel array GPR systems to perform localization. In this evaluation, we demonstrate that a low-cost, off-the-shelf, single-channel GPR system can be used for localization. While our evaluation was performed with a ground-coupled GPR system, our work can generalize for air-launched GPR systems, improving the viability for real world deployment.

Our learned sensor model is trained using data collected at a single point in time and is not currently equipped for long-term prior map registration, since measurements are affected by changes in subsurface moisture content. In the future, we would like to extend this work by using data collected from highly accurate, multi-channel GPR systems during different

seasons and weather conditions to improve the robustness of prior map registration. This is now feasible with the release of the GROUNDED dataset [32].

Current GPR-based localization systems must revisit explored locations to correct for accumulated drift. Future work will investigate methods to attenuate this requirement by modeling the typical relationships of continuous subsurface structures (eg. pipes, subsurface mines, natural caves), enabling more robust GPR-based state estimation in a broader range of settings where revisitation is not practical. We hope to build on this work and move towards general representations for GPR data needed to enable robust, real world deployment of localizing GPR systems.

## VII. ACKNOWLEDGEMENTS

The authors are grateful to Dimitrios Apostolopoulos and Herman Herman for providing the Noggin 500 GPR and Sebastian Scherer for providing the XSENS IMU used to support this study. We thank David Kohanbash for operating the total station used to acquire ground truth positions for this study, Robert Bittner for helping us use the National Robotics Engineering Center facilities for testing, and Ryan Darnley for supporting the initial development of the baseline inertial navigation system.

## REFERENCES

- [1] H. M. Jol, in *Ground Penetrating Radar Theory and Applications*, H. M. Jol, Ed. Amsterdam: Elsevier, 2009.
- [2] L. Pajewski, A. Benedetto, X. Derober, A. Giannopoulos, A. Loizos, G. Manacorda, M. Marciniak, C. Plati, G. Schettini, and I. Trinks, "Applications of ground penetrating radar in civil engineering — COST action TU1208," in *2013 7th International Workshop on Advanced Ground Penetrating Radar*, 2013, pp. 1–6.
- [3] E. G. Ernenwein and K. L. Kvamme, "Data processing issues in large-area GPR surveys: correcting trace misalignments, edge discontinuities and striping," *Archaeological Prospection*, vol. 15, no. 2, pp. 133–149, 2008.
- [4] M. Cornick, J. Koechling, B. Stanley, and B. Zhang, "Localizing Ground Penetrating RADAR: A Step Toward Robust Autonomous Ground Vehicle Localization," *Journal of Field Robotics*, vol. 33, no. 1, pp. 82–102, jan 2016.
- [5] T. Ort, I. Gilitschenski, and D. Rus, "Autonomous navigation in inclement weather based on a localizing ground penetrating radar," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 3267–3274, 2020.
- [6] F. Dellaert and M. Kaess, "Factor graphs for robot perception," *Foundations and Trends in Robotics*, vol. 6, no. 1-2, pp. 1–139, 2017.
- [7] C. Cadena, L. Carlone, H. Carrillo, Y. Latif, D. Scaramuzza, J. Neira, I. Reid, and J. J. Leonard, "Past, present, and future of simultaneous localization and mapping: Toward the robust-perception age," *IEEE Transactions on Robotics*, vol. 32, no. 6, pp. 1309–1332, 2016.
- [8] J. Czarowski, T. Laidlow, R. Clark, and A. J. Davison, "DeepFactors: Real-time probabilistic dense monocular SLAM," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 721–728, 2020.
- [9] P. Sodhi, M. Kaess, M. Mukadam, and S. Anderson, "Learning tactile models for factor graph-based estimation," *arXiv preprint arXiv:2012.03768*, 2020.
- [10] M. Bloesch, J. Czarowski, R. Clark, S. Leutenegger, and A. J. Davison, "CodeSLAM - learning a compact, optimisable representation for dense visual SLAM," *CoRR*, vol. abs/1804.00874, 2018. [Online]. Available: <http://arxiv.org/abs/1804.00874>
- [11] H. Herman, "Robotic Subsurface Mapping Using Ground Penetrating Radar," Tech. Rep., 1997.
- [12] P. Furgale, T. Barfoot, and N. Ghafoor, "Rover-based surface and subsurface modeling for planetary exploration," in *Field and Service Robotics*, A. Howard, K. Iagnemma, and A. Kelly, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2010, pp. 499–508.
- [13] S. Hamran, T. Berger, S. Brovoll, L. Damsgard, O. Hellenen, M. J. Oyan, H. E. Amundsen, L. Carter, R. Ghent, J. Kohler, M. Mellon, D. Paige, D. Plettemeier, and J. Eide, "RIMFAX: A GPR for the Mars 2020 rover mission," in *2015 8th International Workshop on Advanced Ground Penetrating Radar (IWAGPR)*, 2015, pp. 1–4.
- [14] J. Lai, Y. Xu, R. Bugiolacchi, X. Meng, L. Xiao, M. Xie, B. Liu, K. Di, X. Zhang, B. Zhou, S. Shen, and L. Xu, "First look by the Yutu-2 rover at the deep subsurface structure at the lunar farside," *Nature Communications*, vol. 11, no. 1, pp. 1–9, dec 2020.
- [15] K. Nonami, *Development of Mine Detection Robot Mine Hunter Vehicle (MHV), Controlled Metal Detector and Multi-functional Hydraulic Manipulator*. London: Springer London, 2009, pp. 123–143.
- [16] M. Sato, *Principles of Mine Detection by Ground-penetrating Radar*. London: Springer London, 2009, pp. 19–26.
- [17] H. M. La, R. S. Lim, B. Basily, N. Gucunski, J. Yi, A. Maher, F. A. Romero, and H. Parvardeh, "Autonomous robotic system for high-efficiency non-destructive bridge deck inspection and evaluation," in *IEEE Intl. Conf. on Automation Science and Engineering (CASE)*, 2013, pp. 1053–1058.
- [18] R. M. Williams, L. E. Ray, J. H. Lever, and A. M. Burzynski, "Crevasse detection in ice sheets using ground penetrating radar and machine learning," *IEEE J. of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 7, no. 12, pp. 4836–4848, 2014.
- [19] R. Mur-Artal and J. D. Tardós, "ORB-SLAM2: an open-source SLAM system for monocular, stereo and RGB-D cameras," *CoRR*, vol. abs/1610.06475, 2016.
- [20] J. Engel, T. Schöps, and D. Cremers, "LSD-SLAM: Large-scale direct monocular SLAM," in *Eur. Conf. on Computer Vision (ECCV)*, D. Fleet, T. Pajdla, B. Schiele, and T. Tuytelaars, Eds. Cham: Springer International Publishing, 2014, pp. 834–849.
- [21] J. Zhang and S. Singh, "LOAM: Lidar odometry and mapping in real-time," in *Robotics: Science and Systems (RSS)*, 2014.
- [22] P. Fischer, A. Dosovitskiy, E. Ilg, P. Häusser, C. Hazirbas, V. Golkov, P. van der Smagt, D. Cremers, and T. Brox, "FlowNet: Learning optical flow with convolutional networks," *CoRR*, vol. abs/1504.06852, 2015. [Online]. Available: <http://arxiv.org/abs/1504.06852>
- [23] D. Cattaneo, M. Vaghi, A. L. Ballardini, S. Fontana, D. G. Sorrenti, and W. Burgard, "CMRNet: Camera to LiDAR-map registration," *2019 IEEE Intelligent Transportation Systems Conference (ITSC)*, Oct 2019. [Online]. Available: <http://dx.doi.org/10.1109/ITSC.2019.8917470>
- [24] D. Cattaneo, D. G. Sorrenti, and A. Valada, "CMRNet++: Map and camera agnostic monocular visual localization in LiDAR maps," 2020.
- [25] M. Kaess, A. Ranganathan, and F. Dellaert, "iSAM: Incremental smoothing and mapping," *IEEE Trans. Robotics*, vol. 24, no. 6, pp. 1365–1378, 2008.
- [26] M. Kaess, H. Johannsson, R. Roberts, V. Ila, J. Leonard, and F. Dellaert, "iSAM2: Incremental smoothing and mapping using the Bayes tree," *Intl. J. of Robotics Research*, vol. 31, no. 2, pp. 216–235, Feb. 2012.
- [27] X. Feng, M. Sato, Y. Zhang, C. Liu, F. Shi, and Y. Zhao, "CMP antenna array GPR and signal-to-clutter ratio improvement," *IEEE Geoscience and Remote Sensing Letters*, vol. 6, no. 1, pp. 23–27, 2009.
- [28] E. Huber and G. Hans, "RGPR — an open-source package to process and visualize GPR data," in *2018 17th International Conference on Ground Penetrating Radar (GPR)*, 2018, pp. 1–4.
- [29] C. Forster, L. Carlone, F. Dellaert, and D. Scaramuzza, "On-Manifold Preintegration for Real-Time Visual-Inertial Odometry," Tech. Rep. [Online]. Available: <https://youtu.be/CsJkci5lfc0>
- [30] A. Baikovitz, P. Sodhi, M. Dille, and M. Kaess, "Cmu-gpr dataset: Ground penetrating radar dataset for robot localization and mapping," 2021.
- [31] J. Sturm, N. Engelhard, F. Endres, W. Burgard, and D. Cremers, "A benchmark for the evaluation of RGB-D SLAM systems," in *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, Oct. 2012.
- [32] T. Ort, I. Gilitschenski, and D. Rus, "GROUNDED: The Localizing Ground Penetrating Radar Evaluation Dataset," in *Proceedings of Robotics: Science and Systems*, Virtual, July 2021.