Range-based GP Maps: Local Surface Mapping for Mobile Robots using Gaussian Process Regression in Range Space

Margaret Hansen[1*], David Wettergreen[1]

***Abstract*— This work introduces range-based GP maps, which model the range from a LiDAR sensor as a Gaussian process (GP) of the spherical reference frame used to observe the terrain. Such a model aligns the uncertainty in underlying sensor observations with the uncertainty produced from the GP regression. Experimental evaluation on simulated natural terrain indicate that local range-based GP maps perform comparably to elevation-based methods when predicting the local terrain height, with the former producing more stable parameters and a better uncertainty representation. An aggregation method is proposed using the pose as an additional input to the GP. Unlike elevation-based methods, range-based GP maps are capable of modeling overhangs and vertical obstacles with ease, which elevation-based methods are not designed to handle. This is demonstrated through examples of local range-based GP maps built on real-world data from a fully 3D subterranean environment.**

I. Introduction

Terrain mapping is one of the foremost tasks of mobile robots operating in unknown and unstructured environments. Without an accurate model of the surrounding terrain, a mobile robot will not be able to autonomously localize itself and perform path and motion planning. Maps themselves aggregate and reduce dense and memory-intensive perception data while simultaneously preserving necessary features and potentially enriching the model with computed information. Techniques for mapping generally depend on the map's end use and cover a wide variety of methods that can be classified into discrete versus continuous maps, metric versus topological maps, and/or whether the maps have a probabilistic structure (i.e. can model uncertainty about a surface in addition to the surface itself). The latter type of map is of particular interest due to these maps' ability to encode uncertainty information that can be used by the algorithms that rely on them.

The focus of this work is on Gaussian process (GP) regression methods for local surface modeling by a ground-based mobile robot. GP regression models are useful for a number of reasons, including their modeling of spatial correlations and uncertainty estimates and the potential to query them at any location. Gaussian process-based mapping has focused on elevation, which provides uncertainty in elevation only. For aerial robots, this is reasonable since they observe the terrain from a "top-down" view, which more closely aligns uncertainty in elevation with the range uncertainty from the sensor. However, as shown in Figure 1, ground-based robots observe the terrain from an oblique view, meaning the range uncertainty from the sensor is not aligned with the z component of the measured terrain points. For the latter, a range-based approach is better suited to correctly capture the sensor uncertainty in the terrain model.

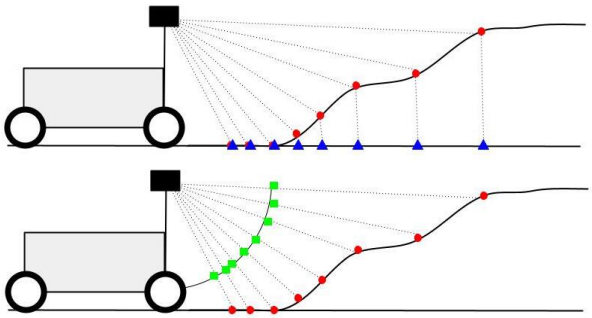

(a) Ground-based range measurements.

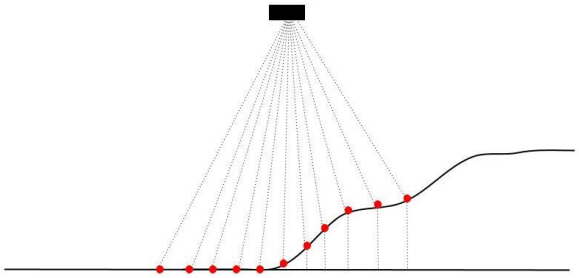

(b) Aerial range measurements.

Fig. 1: Ground-based and aerial robots observe terrain from different perspectives. Using angular coordinates along a sphere to parameterize the range values (a, bottom) aligns predicted uncertainty with the angle of incidence of a LiDAR beam, which is not the case when using 2D ground coordinates to parameterize the height (a, top). Aerial measurements (b) show less pronounced difference between the two methods due to the top-down viewing angle.

We present range-based GP maps, maps built around the range sampling structure of a LiDAR and represented by a GP. Rather than model elevation as a function of 2D position, these maps model the sensor output (range) directly as a function of the inclination and azimuth angles used to take measurements. Doing so allows the local mapping method to model uncertainty in the range measurement itself, which is then translatable into its 3D components based on knowledge of the sensor position and sampling angles. Additionally, these range-based maps implicitly encode visibility infor-

[1] Robotics Institute, Carnegie Mellon University, Pittsburgh, PA, USA.
[*] Corresponding author. margareh@cs.cmu.edu.

mation into the map and enable modeling of overhanging structures. The implicit limitation on uncertainty in elevation provided by visibility information in range-based GP maps is shown in the simple example in Figure 2. In summary, range-based GP maps provide all of the advantages of an elevation-based Gaussian process map but with three additional benefits: (1) more correct modeling of sensor uncertainty for ground-based robots; (2) modeling of overhangs or vertical objects and structures; and (3) implicit encoding of visibility information.

II. Related Work

Gaussian process regression and classification have both been widely considered in the robotics community for the task of mapping an environment. The appeal of GP-based methods is their ability to incorporate spatial correlation in a continuous manner, enabling interpolation and extrapolation for query points either between or outside of sample locations. GP methods have been used to model surface geometry itself by expanding on basic elevation models for both local and large-scale mapping [1,2], as well as integrating Gaussian process modeling into other map representations such as occupancy grids [3] and signed distance fields [4].

A. Large-Scale Elevation Mapping

Elevation modeling using Gaussian process regression has a long history in the geostatistics community, here termed "kriging." These models typically involve using a GP in concert with a digital elevation model (DEM) where GP regression is used to fill in portions of the DEM where no observations were taken [5]. Some robotic mapping methods have incorporated elevation-based GP mapping in ways similar to kriging, including methods that divide the environment into clusters and generate individual Gaussian process regressions for each cluster patch [2]. GPGM-SLAM uses Gaussian process models to interpolate elevation across local submaps, which are then used to detect potential loop closures [6]. While the method itself shows promise as a mapping method that incorporates uncertainty, this uncertainty is not ultimately used in the map's application (finding loop closures).

B. Local Elevation-based Terrain Modeling

Elevation-based GP regression has also been used on the local scale in mobile robotics, where the $O(n^3)$ computational cost is not as severe. To handle discontinuities, non-stationarity can be incorporated into the kernel based on local covariance matrices for individual data points [1]. These covariance matrices can be adapted over time to "learn" the structure of the dataset, including the presence of discontinuities. Visibility information can also be encoded via a kernel regression-based framework similar to a GP [7]. This visibility constraint is important but used in an optimization of the weights on individual data points, meaning the kernel function itself is unaffected (and indeed not directly computed with this method).

C. Modeling Signed Distance Fields and Range

Gaussian process methods have also been used to predict signed distance fields (SDFs). These models are more similar to the range-based GP method proposed here. Termed "Gaussian Process Implicit Surfaces (GPIS)," these model the range based on angle in the 2D context for estimating a signed distance field; the surface model then becomes the zero crossings of this SDF [4]. GPIS have been extended to include an ensemble of GPs rather than one model, as well as to perform approximations based on independence assumptions and intelligent data limitations (selection of inducing points) to improve runtime [8]. While this is certainly an interesting environment model, using the SDF for surface modeling requires a raytracing operation to find the zero crossings that is more computationally expensive than direct prediction from the GP. Finally, and most similar to the current work, GPs have been used in polar coordinates for interpolation between range measurements for densification of LiDAR scans [9]. However, no terrain modeling was performed with this GP.

While elevation-based GP methods have historically been popular due to their ease of use and long history of use in geostatistics, such models do not account for the oblique view of a ground-based robot. As a result, visibility constraints must be added on externally, removing the ability to query uncertainty directly from the model. Recent work using SDFs is more in line with the oblique viewpoint; however, these methods focus on 2D indoor environments when performing tests for mobile robots.

III. Methodology

A. Local Range-based GP Maps

Range measurements taken from a 3D LiDAR scanner are indexed based on their azimuth (ϕ) and inclination (θ), with the measured range r specifying the third coordinate necessary to fix a point in 3D space. With $\vec{\alpha} = (\theta, \phi)^T$ as the vector of the two angles, the range can be modeled as

$$R(\vec{\alpha}) = r(\vec{\alpha}) + \epsilon_R \tag{1}$$

$r(\vec{\alpha})$ is the range to the surface for a given set of angular coordinates $\vec{\alpha}$ and $\epsilon_R(\vec{\alpha})$ is the error associated with a given measurement. Gaussian process regression can be used to model the range by assuming

$$r(\vec{\alpha}) \sim GP(m(\vec{\alpha}), K(\vec{\alpha}, \vec{\alpha}')) \tag{2}$$

which is a Gaussian process model with mean function $m(\vec{\alpha})$ and kernel function $K(\vec{\alpha}, \vec{\alpha}')$. This type of Gaussian process regression aligns the kernel uncertainties with the rays between the LiDAR sensor and the surface intersection points being measured. Since azimuth, inclination, and range are the coordinates being used, this is akin to modeling range in a spherical space as is shown in Figure 1(a). While an elevation-based GP would project the elevation value onto the xy-plane, the range-based GP projects range measurements onto a sphere of fixed radius. The projection vectors of the former do not align with the sensor viewing

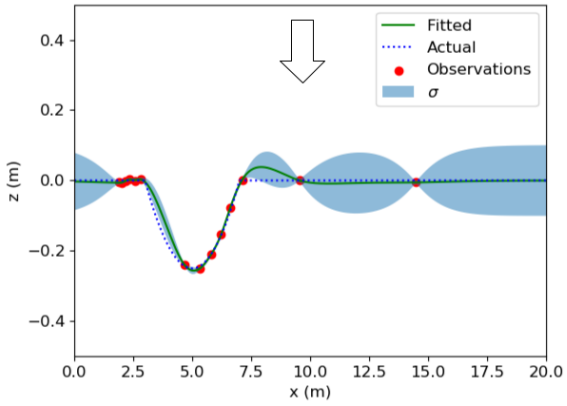

(a) Elevation-based model using a Matérn kernel with hyperparameters ν = 1.5, κ = 1.76, σ = 0.0992, and η = 0.00001 and a zero elevation prior.

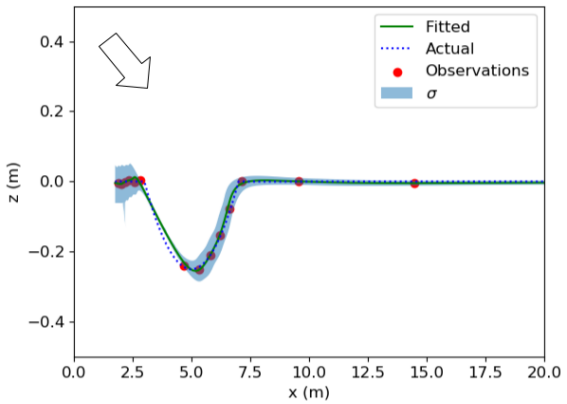

(b) Range-based model using a Matérn kernel with hyperparameters ν = 1.5, κ = 0.0445, σ = 0.813, and η = 0.0137 and a flat ground prior.

Fig. 2: Simple example of elevation versus range-based GP modeling in 2D. Arrows indicate the viewing angle associated with the model: "top-down" for the elevation GP and oblique for the range GP. The σ displayed in the plot is the predicted standard deviation of z; for the range-based GP, some of the uncertainty in the range measurements is translated into the x direction, while for the elevation-based GP all uncertainty is in the z direction. Also note the presence of a small incorrect positive "obstacle" at around 7.5 m in the elevation-based GP model that is not present in the range-based model.

angles, as is the case for the latter; as a result, the elevation-based GP will not model the sensor uncertainty correctly but the range-based GP is expected to do so.

The kernel function $K(\vec{\alpha}, \vec{\alpha}')$ can be any of the typical covariance kernels used for Gaussian process regression, such as squared exponential or rational quadratic. However, the geodesic (angular) distance does not necessarily produce valid kernels, so Euclidean (or chordal) distance between points in 3D space is used instead to ensure kernel validity [10]. Euclidean coordinates produced when using the maximum range R_{max} in the conversion from spherical coordinates will be used in these computations:

$$\tilde{x} = R_{max} \sin\theta \cos\phi \tag{3}$$

$$\tilde{y} = R_{max} \sin\theta \sin\phi \tag{4}$$

$$\tilde{z} = R_{max} \cos\theta \tag{5}$$

This is equivalent to computing the Euclidean coordinates on the "range sphere" where the range equals R_{max} always, which produces a valid kernel for Gaussian process modeling on a sphere [10]. For this work, a Matérn kernel with ν fixed to 1.5 was chosen to better model discontinuities:

$$K(\vec{\tilde{x}}, \vec{\tilde{x}}') = \sigma^2 \frac{1}{\Gamma(\nu)2^{\nu-1}} \left(2\sqrt{\nu}\frac{d}{\kappa}\right)^{\nu} K_\nu \left(2\sqrt{\nu}\frac{d}{\kappa}\right) \tag{6}$$

d is Euclidean distance between $\vec{\tilde{x}}$ and $\vec{\tilde{x}}'$, ν is the differentiability parameter (set to 1.5), κ is the length scale, and $K_\nu(\cdot)$ is the modified Bessel function of the second kind [11]. For this work, the Matérn kernel is additively combined with a White kernel (parameterized by the noise level η) to account for underlying noise in the data.

It is typical when training a Gaussian process regression to use a mean function of zero; however, for this application, doing so would mean using a prior of zero range, which does not align with any reasonable expectation of range observations. Two options for mean functions are available in this case: a sphere of pre-determined non-zero radius (e.g. constant range) and a flat plane at ground height, transformed into the expected range measurements for a given set of inclination and azimuth. The latter is computed as:

$$m(\vec{\alpha}) = \frac{s_z}{\sin(\theta)\cos(\phi)\sin(\beta) - \cos(\theta)\cos(\beta)} \tag{7}$$

In the equation above, s_z is the (assumed known) sensor height relative to the ground and β is the sensor pitch with respect to the horizontal. This mean function produces the predicted range for given azimuth and inclination angles along a flat ground plane and is derived from the simpler equation

$$\hat{r} = \frac{s_z}{\sin(\tilde{\theta} - 90°)} \tag{8}$$

in which $\tilde{\theta}$ is the inclination angle after adjusting for the sensor pitch. Equation 7 can be derived based on the Euclidean and spherical coordinate transforms between the sensor frame and an equivalently centered frame that is not pitched relative to the ground (i.e. the frame in which $\tilde{\theta}$ is measured).

It is anticipated that the flat ground plane will be of use for environments with little to no vertical or overhanging features. The spherical prior is anticipated to be useful in cases with many overhanging or vertical features, such as a robot operating underground or in environments with vertical relief like near the bases of cliffs.

Given a training dataset $\vec{\alpha}_{tr}$ with observed range values $\vec{r}_{tr}$, inference is performed for a set of elevation and azimuth angles $\vec{\alpha}_* = (\theta_*, \phi_*)$ using the usual inference equations for a Gaussian process [12]:

$$\vec{r}_* = m(\vec{\alpha}_*) + K(\vec{\alpha}_*, \vec{\alpha}_{tr})^T \Lambda \, \vec{r}_{tr} \tag{9}$$

$$\sigma^2_{r_*} = K(\vec{\alpha}_*, \vec{\alpha}_*) - K(\vec{\alpha}_*, \vec{\alpha}_{tr})^T \Lambda \, K(\vec{\alpha}_*, \vec{\alpha}_{tr}) \tag{10}$$

where $\Lambda = (K(\alpha_{tr}, \alpha_{tr}) + \sigma^2 I)^{-1}$ is the inverse kernel matrix (in this work, the addition of the noise is contained in the White kernel). For a given set of kernel hyperparameters, training the GP is equivalent to computing this inverse.

B. Aggregating the Map

One drawback to range-based GP mapping as outlined above is that it is an inherently local model. The sensor measurement angles are defined with respect to the sensor position - if the sensor moves, the same sampling angles will not correspond to the same points in the world. Additionally, modeling range as a function of angles means that one viewing direction may only have one range value for a given local map. However, as the robot moves, the map should be capable of being updated and accumulated in order to more fully model the environment - indeed, Gaussian processes themselves are easily updated with incoming data. In this work, the model is updated by simply incorporating the pose as additional input variables to the GP regression such that $\vec{\alpha}_{world} = (\theta, \phi, x, y, z, \gamma)^T$ is the new set of inputs, with x, y, and z representing the translational components of the sensor pose and γ the sensor heading. One practical implication of including the pose is that two scans are necessary before modeling can begin in the aggregate. The range prior is still computed based solely on the observation angle parameters, and the kernel computations are updated to use Euclidean distance for the translational components of the pose and chordal distance for the observation angles and heading. Only the heading and translational components of the pose are currently integrated into the GP model.

IV. Experiments

Range-based GP mapping has been tested with data representing a ground robot operating on the surface of simulated cratered terrain, with no vertical obstacles or overhangs. Two sets of experiments were run: local mapping and aggregate mapping. The local mapping experiments test the range-based GP method's ability to represent the environment using a single LiDAR scan, while the aggregate mapping experiments test its ability to represent the environment using multiple LiDAR scans and known robot pose.

The terrain model, shown in Figure 3 with an example robot trajectory, represents a 100 m x 100 m heightmap with a resolution of 5 cm and is generated by sampling from the size-frequency distribution for lunar craters described in NASA's Design Specification for Natural Environments [13], extrapolated to represent smaller craters. The crater models themselves are built based on a parabolic estimation of crater shape as a function of diameter [14].

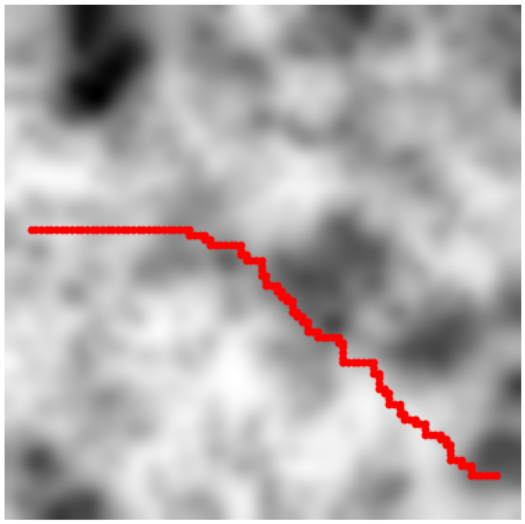

Fig. 3: Cratered terrain heightmap with example trajectory used for simulated experiments with a planetary rover. Lighter regions are higher elevation and darker regions are lower elevation.

Sensor observations are generated at a height of 75 cm from the ground at sample points across the terrain model. These observations are generated via raycasting from the sensor position at elevations ranging from -15 degrees to 0 degrees (relative to the horizontal) with a resolution of 1 degree and azimuths ranging from -90 degrees to 90 degrees with a resolution of 2 degrees. Raycasting was performed out to 10 m, and points at the maximum range were removed from the scans prior to processing. 20% of observations from each simulated scan are used for training and have noise added to the range measurements; the other 80% was reserved for testing to limit the amount of training data provided to the GP. The simulated noise added to the training data is scaled linearly by the range value such that at 29 m, the standard deviation of the noise is 0.75 cm in accordance with the noise properties of a Velodyne VLP-16 sensor [15].

For local mapping experiments, random sample points are enough. However, for the aggregate mapping experiments, a simple trajectory is generated using A^* between randomly sampled start and end locations with a threshold on the absolute height difference from the mean to define obstacles. This trajectory is downsampled and the heading computed based on the direction of travel to produce the positions at which raycasting is performed for these experiments.

The White kernel noise level is fixed to 0.0075 (0.75 cm) for all tests. The other kernel hyperparameters are selected by running scikit-learn's built-in optimizer [16] on the local range-based GP model over 100 randomly-selected scans and averaging the resulting values; these parameters were

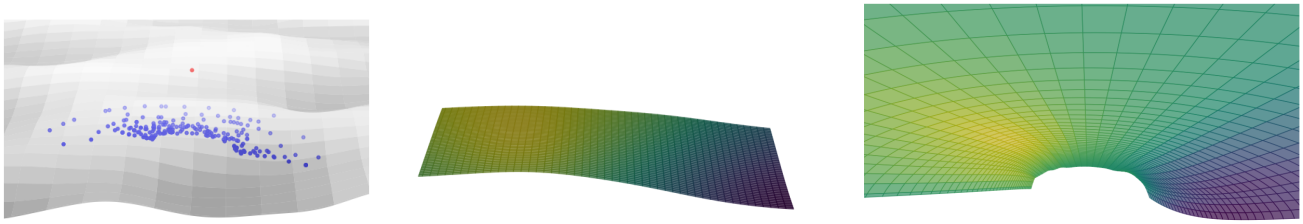

(a) Robot position (red) and raytraced scan (blue).

(b) Elevation-based GP model.

(c) Range-based GP model.

Fig. 4: Simulated scan from cratered terrain (left), along with the local map produced using an elevation-based GP (center) and a range-based GP (right). Viewing directions for all images are in line with the sensor's viewpoint. Both GPs are colored based on elevation (yellow is higher and purple is lower). Both methods are capable of modeling the negative obstacle to the right of the scan, with the range-based GP returning to the flat prior more quickly than the elevation-based GP.

initialized to $\sigma = 2.5$ and $\kappa = 0.15$ based on initial tests. The performance of local range-based GP mapping was then assessed over 100 separate randomly selected scans and compared to an elevation-based GP regression map, using the root mean squared error (RMSE) in z over the test data and the average uncertainty from the GP models. The performance of the aggregate range-based GP mapping method was tested on 5 trajectories across the terrain map, from which scans were aggregated every 5 steps in the trajectory (representing every 5 meters).

V. Results & Discussion

A. Local Mapping

Local mapping tests indicate that the range-based GP maps are capable of representing local terrain based on a single scan. An example comparison between the elevation and range-based GP maps can be seen in Figure 4, along with the underlying surface and raytraced sample points. Both the elevation and range-based GP maps are capable of detecting and modeling the negative obstacle to the right of the scan. The performance statistics in Table I support this observation; the RMSE in z is comparable between the range-based GP and the elevation-based GP. The uncertainty in z is lower for the range-based model, which aligns with expectations based on the oblique view used in this model - some of the range uncertainty is translated into the x and y components of the ray along which the observation is taken. Unlike the elevation-based GP, the range-based GP is not explicitly designed to model elevation, yet it is still capable of locally predicting elevation on par with the elevation-based GP.

The average optimal hyperparameters from the 100 optimization scans are also shown in Table I, along with their standard deviations (in parentheses). The length scale hyperparameter for the range-based GP is much more stable than that for the elevation-based GP, with the latter having a much larger standard deviation and a more spread out distribution shown in Figure 5. 10 optimization scans produced a length scale of 50 for the elevation GP, which was the maximum allowed length scale for both models; in these cases, the optimizer may not have found what it considered to be

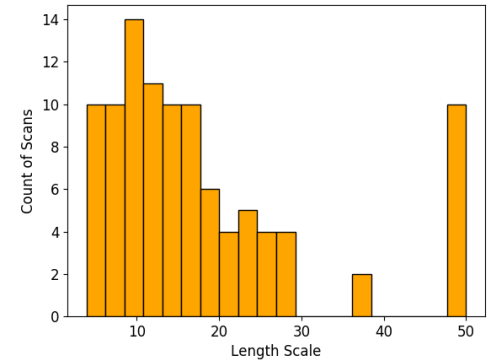

(a) Elevation-based GP model.

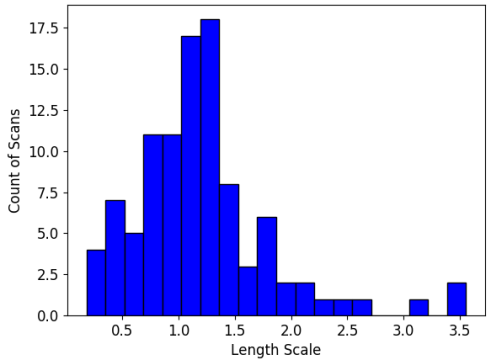

(b) Range-based GP model.

Fig. 5: Histogram of optimized length scale parameter for elevation-based and range-based GP models, across 100 training scans from the simulated cratered terrain dataset. Note the large spread of length scale values for the elevation-based GP model, including 10 scans where the length scale selected was 50 (the upper bound). The range GP models have a much tighter distribution of optimal length scale values.

"optimal" due to the bound. This indicates that the range-based GP model will likely be more generalizable to test data once a length scale has been chosen for use.

	Range-based GP	**Elevation-based GP**
Kernel scalar	2.302 (0.980)	0.0842 (0.0497)
Length scale	1.206 (0.608)	18.102 (12.682)
Local RMSE	0.0121 (0.0048)	0.0085 (0.0021)
Local GP uncertainty		
x	0.0603 (0.0009)	–
y	0.0620 (0.0011)	–
z	0.0169 (0.0004)	0.0890 (0.0003)

TABLE I: Optimized hyperparameters and performance metrics for range-based GP maps compared to elevation-based GP maps generated from 100 test scans randomly sampled from the cratered terrain heightmap. RMSE is computed for z (elevation) over held-out test points. Quantities in parentheses are standard deviations of the performance metrics.

Range-based GP maps are comparable in terms of predicting elevation to elevation-based GP maps, have a more stable length scale parameter, and better represent the underlying sensor uncertainty. As a result, range-based GP maps should be preferred to elevation-based methods for applications using ground robots that require a good model of uncertainty.

B. Aggregate Mapping

The aggregate mapping tests indicate that the range-based GP model does not predict elevation as well when operating in this mode, though it is still capable of modeling general trends in the data. The full model after processing the trajectory from Figure 3 is shown in Figure 6(a). The shading in this model corresponds to the height for ease of comparison to the heightmap; the aggregate model exhibits many of the same terrain features as the heightmap, such as the darker regions in the middle of and towards the end of the trajectory and the lighter region between these. The absolute error between the height predicted from the range-based GP and the actual heightmap, shown in Figure 6(b), indicates that the model is particularly good at representing regions directly underneath the robot (along the trajectory), as well as flat regions, likely due to the flat prior. The regions with the largest error tend to be towards the edges of the scans or correspond to regions with more abrupt changes.

The aggregate range-based GP maps do not perform as well as a batch elevation-based GP method at predicting height over the test scans across the 5 trajectories tested: the average RMSE in elevation at the end of the range-based GP aggregation is 0.1442, while the RMSE of the elevation-based model is 0.0096. However, the elevation-based GP model is run in batch and does not integrate the pose in the same manner as the range-based GP model, meaning the two models are not quite equal for the sake of comparison. Again, the elevation-based GP model is designed to explicitly model elevation, while the range-based model is not. Despite its worse performance, the range-based model shows promise for aggregate mapping if improved upon.

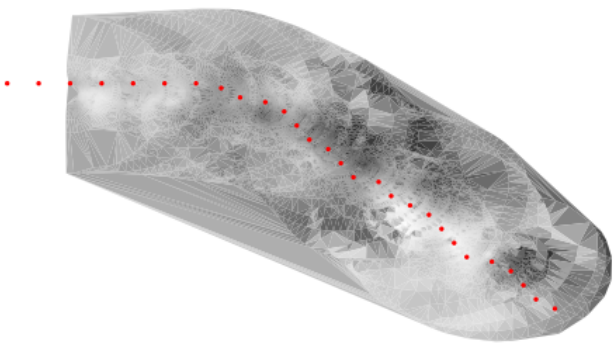

(a) Aggregate range-based GP map.

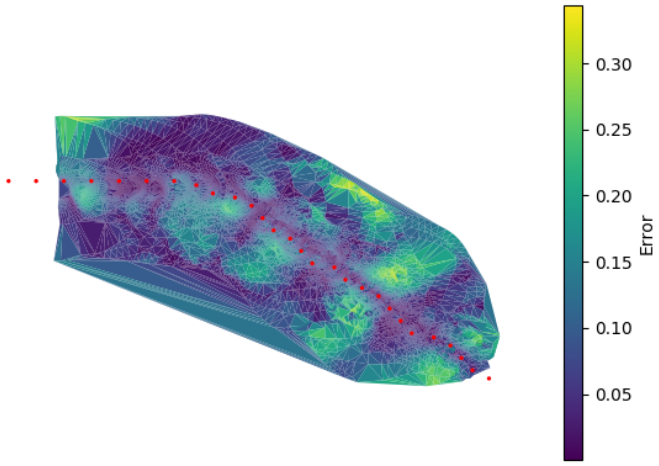

(b) Absolute error between predicted height from range-based GP map and heightmap ground truth.

Fig. 6: Aggregated range-based GP map produced using trajectory shown in Figure 3, downsampled to one scan every 5 meters. Red points are trajectory points used to build the map. The elevation (top) roughly corresponds to the elevation shown in the heightmap; the error (bottom) is lower near the trajectory and higher for more pronounced terrain features. Points near the end of the trajectory were extrapolated off of the heightmap and have no associated ground truth or computed error.

C. Real-world Data with Overhangs

Finally, to demonstrate the range-based GP maps' capabilities of modeling vertical and overhanging structures, local maps were generated from real-world data collected from a ground robot operating in a subterranean environment with a Velodyne VLP-16 LiDAR sensor. This data was collected during the final DARPA SubT challenge by CMU's Team Explorer [17]. 50 local scans were used to compute the optimal hyperparameters, which were $\sigma = 5.16$ and $\kappa = 0.177$. Again, the noise level was fixed to 0.0075. For these tests, a spherical prior was used with the maximum range set to 10 instead of the flat ground prior since the environment contains many vertical and overhanging features. An example local scan and associated local range GP map are shown in Figure 7 from two angles. Note how the method is capable

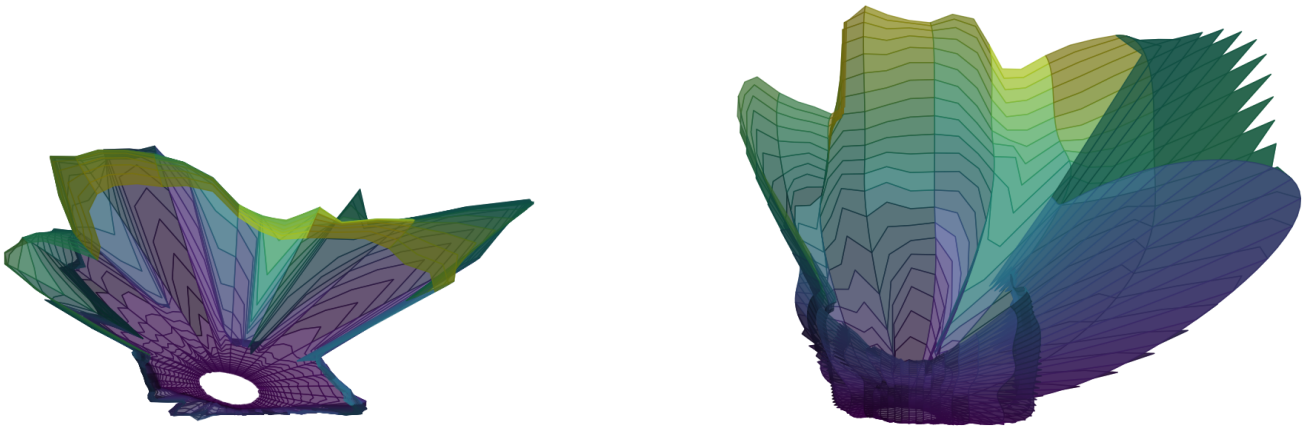

(a) Top down view.

(b) Front view.

Fig. 7: Example LiDAR scan in a subterranean environment with fitted range GP model. Robot is positioned at the center of the circle shown in the top-down view (left) and observes a few walls and an open area. Color indicates z value (yellow is higher and purple is lower). Note the modeling of the vertical wall at the bottom of the top-down view and at the front of the front view.

of modeling the vertical walls at the bottom and sides of Figure 7(a), as well as the large open area to the top of this figure. An elevation-based GP map would not be capable of modeling such features because the underlying scan would produce many overlapping points on the ground plane.

VI. Conclusion

This work proposes range-based Gaussian process regression for modeling unstructured terrain in 3D based on a LiDAR sensor. These models better align with the oblique viewing angle used when observing terrain from a ground-based robot than elevation-based methods, resulting in local maps that better represent sensor uncertainty, are capable of modeling overhangs and vertical structures, and implicitly encode the sensor visibility. Tests on simulated cratered terrain indicate that range-based GP maps are capable of locally producing estimates with comparable elevation error to an elevation-based model, though with uncertainty of the range measurement itself modeled. Additionally, aggregation across maps produces results that are at a high level similar to the traversed terrain, though extrapolations beyond the extent of the model are worse for the range-based GP than when modeling elevation. Initial results from local range-based GP maps built from subterranean LiDAR scans indicate that this method is capable of modeling overhangs and vertical structures, unlike an elevation-based method.

There are many aspects of range-based GP mapping that have not been empirically analyzed or focused on in this paper but are still important aspects of the modeling procedure. In particular, no steps have been taken to reduce the runtime and memory requirements of the algorithm, both of which are non-trivial. Additionally, only one built-in optimizer has been used to perform hyperparameter optimization, and the range prior was chosen *a priori* and is not algorithmically determined based on the environment. Additionally, the range-based GP models contained in this work are simple and do not account for non-stationarity in the inputs provided to the GP regression.

Future work will investigate how to better model discontinuities and allow different hyperparameters for different input variables. Further testing on real-world data from various scenarios (subterranean and other) will be performed to fully validate the method. Additionally, a number of improvements can be made to decrease runtime and allow the algorithm to adapt to its surroundings - for instance, selecting inducing points that provide the most information to the GP model instead of randomly sampling, and efficiently adapting hyperparameters and the model prior based on the changing terrain as the robot moves. Finally, applications of these maps will be tested, such as using optimization of the pose inputs in the aggregate GP [18] to perform localization. The full benefits of range-based GP mapping are expected become apparent when performing localization or when using the uncertainty in planning algorithms due to the better representation of uncertainty incorporated in these models.

Acknowledgements This work was supported by NASA Space Technology Graduate Research Opportunity grant 80NSSC22K1206 and NASA Solar System Exploration Virtual Institute and the Planetary Science Institute (SSERVI) project TREX: Toolbox for Research and Exploration (80ARC017M0005).

References

[1] T. Lang, C. Plagemann, and W. Burgard, "Adaptive Non-Stationary Kernel Regression for Terrain Modeling," in *Robotics: Science and Systems III*. Robotics: Science and Systems Foundation, Jun. 2007.

[2] X. Liu, D. Li, Y. He, and F. Gu, "Efficient and Multifidelity Terrain Modeling for 3D Large-Scale and Unstructured Environments," *Journal of Field Robotics*, pp. 1–37, 2022.

[3] M. Ghaffari Jadidi, J. Valls Miro, and G. Dissanayake, "Gaussian processes autonomous mapping and exploration for range-sensing mobile robots," *Autonomous Robots*, vol. 42, no. 2, pp. 273–290, Feb. 2018. [Online]. Available: http://link.springer.com/10.1007/s10514-017-9668-3

[4] B. Lee, C. Zhang, Z. Huang, and D. D. Lee, "Online Continuous Mapping using Gaussian Process Implicit Surfaces," in *2019 International Conference on Robotics and Automation (ICRA)*, May 2019, pp. 6884–6890, iSSN: 2577-087X.

[5] M. Habib, "Evaluation of DEM interpolation techniques for characterizing terrain roughness," *CATENA*, vol. 198, p. 105072, Mar. 2021. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0341816220306226

[6] R. Giubilato, C. L. Gentil, M. Vayugundla, M. J. Schuster, T. Vidal-Calleja, and R. Triebel, "GPGM-SLAM: a Robust SLAM System for Unstructured Planetary Environments with Gaussian Process Gradient Maps," pp. 1721–1753, Aug. 2022. [Online]. Available: https://doi.org/10.55417/fr.2022053

[7] R. Hadsell, J. A. Bagnell, D. Huber, and M. Hebert, "Accurate Rough Terrain Estimation with Space-Carving Kernels," in *Robotics: Science and Systems V*. Robotics: Science and Systems Foundation, 2009, p. 8. [Online]. Available: http://www.roboticsproceedings.org/rss05/p19.pdf

[8] J. A. Stork and T. Stoyanov, "Ensemble of Sparse Gaussian Process Experts for Implicit Surface Mapping with Streaming Data," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*, May 2020, pp. 10 758–10 764, iSSN: 2577-087X.

[9] W. Chen, J. Sun, Z. Zhao, and Q. Zheng, "Gaussian Processes in Polar Coordinates for Mobile Robot Using SE(2)-3D Constraints," *Journal of Intelligent & Robotic Systems*, vol. 103, no. 4, p. 72, Dec. 2021. [Online]. Available: https://link.springer.com/10.1007/s10846-021-01520-0

[10] J. Guinness and M. Fuentes, "Isotropic covariance functions on spheres: Some properties and modeling considerations," *Journal of Multivariate Analysis*, vol. 143, pp. 143–152, Jan. 2016. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/S0047259X15002109

[11] C. J. Paciorek and M. J. Schervish, "Nonstationary Covariance Functions for Gaussian Process Regression."

[12] C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning*. MIT Press, 2006.

[13] "Cross-program design specification for natural environments (dsne)," NASA, 2021, sLS-SPEC-159 Revision I, NASA. Accessed at https://ntrs.nasa.gov/citations/20210024522. [Online]. Available: https://ntrs.nasa.gov/citations/20210024522

[14] J. E. Chappelow, "On the depths and shapes of the freshest kilometer-scale simple craters on the lunar maria: A new crater shape model." *Meteoritics & Planetary Science*, vol. 53, no. 4, pp. 813–825, 2018.

[15] J. R. Kidd, "Performance evaluation of the velodyne vlp-16 system for surface feature surveying," Master's thesis, University of New Hampshire, Durham, 2017.

[16] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and Édouard Duchesnay, "Scikit-learn: Machine learning in python," *Journal of Macine Learning Research*, vol. 12, no. 85, pp. 2825–2830, 2011.

[17] Carnegie Mellon University AirLab, "SLAM Datasets: Multi-Robot, Multi-Degraded SLAM Datasets." [Online]. Available: https://superodometry.com/datasets

[18] A. Wiens, W. Kleiber, K. R. Barnhart, and D. Sain, "Surface Estimation for Multiple Misaligned Point Sets," *Mathematical Geosciences*, vol. 52, no. 4, pp. 527–542, May 2020. [Online]. Available: http://link.springer.com/10.1007/s11004-019-09802-y