

Continuum Robot Localization using Distributed Time-of-Flight Sensors

Spencer Teetaert^{*}, Giammarco Caroleo[†], Marco Pontin[†], Sven Lilge[‡],
Jessica Burgner-Kahrs^{*}, Timothy D. Barfoot^{*}, Perla Maiolino[†]

^{*}University of Toronto Robotics Institute, Canada

[†]Oxford Robotics Institute, University of Oxford, UK

[‡]Toronto Metropolitan University, Canada

Abstract—Localization and mapping of an environment are crucial tasks for any robot operating in unstructured environments. Time-of-flight (ToF) sensors (e.g., lidar) have proven useful in mobile robotics, where high-resolution sensors can be used for simultaneous localization and mapping. In soft and continuum robotics, however, these high-resolution sensors are too large for practical use. This, combined with the deformable nature of such robots, has resulted in continuum robot (CR) localization and mapping in unstructured environments being a largely untouched area. In this work, we present a localization technique for CRs that relies on small, low-resolution ToF sensors distributed along the length of the robot. By fusing measurement information with a robot shape prior, we show that accurate localization is possible despite each sensor experiencing frequent degenerate scenarios. We achieve an average localization error of 2.5 cm in position and 7.2° in rotation across all experimental conditions with a 53 cm long robot. We demonstrate that the results are repeated across multiple environments, in both simulation and real-world experiments, and study robustness in the estimation to deviations in the prior map.

I. INTRODUCTION

Continuum robots (CRs) are well-suited for many inspection and manipulation tasks in constrained environments due to their inherent compliance and capacity to be miniaturized [26]. This has led them to be studied for applications such as surgery [5, 13], aircraft inspection [10, 12, 37], and search and rescue operations [39]. This miniaturization comes at the cost of additional size constraints on the sensors that can be integrated into the robot, making localization using CRs particularly challenging. Time-of-flight (ToF) sensors have been studied for CRs recently [1, 7, 18] though the sensor resolution, range, and field of view are limited at the smaller scales of CRs. This can lead to degenerate measurements where single ToF sensors are unable to observe enough features in the environment to fully constrain their positions.

In previous works, distributed sensing has been shown to be an effective way to both localize and perform shape estimation for robots with an estimated configuration [18, 30, 32, 39]. By providing multiple sensing locations along the body of the robot, distributed sensing can provide a richer understanding of both the robot’s configuration and its surrounding environment. Additionally, recent works have demonstrated the effectiveness of continuous-time state estimation for CRs [32, 33], allowing for state estimation from both low- and high-frequency, asynchronous sensor measurements. This



Fig. 1. A continuum robot equipped with distributed sensors is shown localizing in a cluttered environment. The 53 cm long, 7.6 cm diameter extensible soft, continuum robot is equipped with three distributed sensor rings, each with three low-resolution ToF sensors and one gyroscope. Using only these sparse measurements, the robot is able to localize itself with high accuracy (1.9 cm average distance error in the shown environment). Noisy ToF sensor measurements are shown as green points alongside the ToF sensor fields of view. At many points in time, individual sensor rings do not see enough features to fully constrain their location, yet by fusing the distributed measurements with a robot shape prior, accurate localization is still possible.

framework is particularly well-suited for CRs as it enables the fusion of measurements from multiple distributed sensors along the robot’s body with varying sampling rates.

Despite recent progress in distributed sensing and continuous-time state estimation for continuum robots, most existing localization approaches rely on either dense sensing, rigid-body assumptions, or external reference systems. These assumptions break down for small-diameter, deformable robots operating in cluttered, unstructured environments, where on-board sensors are necessarily sparse, low-resolution, and frequently produce geometrically degenerate measurements. In such settings, individual sensors are insufficient to fully constrain robot pose, and even moderate deviations between the prior map and the true environment can lead to estimator failure. In this work, we address this gap by demonstrating that full-body, onboard-only localization of a deformable

continuum robot is feasible using sparse, distributed time-of-flight sensing, even in the presence of frequent measurement degeneracies and partial map mismatch. The contributions of this work are:

- A continuous-time, factor-graph-based MAP estimation framework that integrates sparse time-of-flight measurements with a robot shape prior, enabling stable full-body localization of deformable continuum robots under frequent geometric degeneracies and partial prior map mismatch.
- The first real-world demonstration of onboard-only localization of a soft, extensible continuum robot in cluttered environments using distributed, low-resolution ToF and inertial sensing, validated across both simulated and physical platforms.

The code used in this work is available to the community at https://github.com/utiasASRL/space_time_continuum.

II. RELATED WORK

Continuum and snake-like robots have been proposed for a number of inspection and maintenance tasks in constrained environments. Their ability to be miniaturized [40] enables applications in confined and convoluted environments. Both [10] and [37] explore the use of CRs for aircraft inspection, while [12] demonstrates basic machining tasks in a similar environment. This application space has even seen some commercial success with articulated borescopes [2]. While many previous works have focused on shape estimation of CRs [14, 21, 25, 28], we focus rather on the task of localizing the robot with respect to some inertial reference frame. To enable any degree of autonomy in the aforementioned applications, accurate localization with respect to an environment is critical.

Localization of CRs has been explored through a variety of sensing modalities, including cameras, ToF sensors, and electromagnetic (EM) trackers [29]. Many existing works use sensors with off-robot components for localization, including ultrasound [20], electromagnetic pose sensors [15], fluoroscopic-like stereo imaging [4], and external cameras [24, 27]. These approaches often require line-of-sight to the robot or specialized imaging equipment to sense the robot’s state, which may not be available in many applications. One alternative is to mount the sensors directly on the robot itself and use them to sense the robot’s environment, inferring location from relative positioning. Such approaches are typically more challenging due to the size and weight constraints of the sensors.

Some works have shown promise with camera-based localization [16, 30, 39]. In [35], cameras are used alongside a ToF sensor (lidar) for localization of a snake-like robot. A tip camera as well as a fiber Bragg grating (FBG) sensor is used in [36] to estimate robot pose. In [30], multiple cameras are distributed along a hybrid soft-rigid robot to perform visual SLAM. Through sensor distribution, the authors are able to achieve full-state localization and configuration estimation. As for sensor distribution, [22] proposes an uncertainty-based framework for optimizing sensor placement along CRs

while [19] explores optimal placement of fiber Bragg gratings for shape sensing. Both of these works highlight the importance of sensor placement for accurate state estimation beyond just at the robot tip, though the latter only measures intrinsic information about the robot’s shape and not the location of the robot within a broader environment.

ToF sensing offers different advantages to cameras, enabling native depth perception and increased robustness in visually degraded environments. At smaller scales, however, dense ToF measurements are typically not available. Some existing works have explored the use of sparse ToF sensors in CR applications. In [18], a series of 40-160 small ToF sensors are used for full-body localization of a rigid snake robot. The authors solve for spatial states in series along the robot body in a simulated environment. In [8], a characterization of the VL53L5CX ToF sensor is provided, while [7] proposes a state estimation algorithm using the sensor on a soft robot. In [1], multiple single-beam ToF sensors (VL6180X) are integrated into a proposed sensor module for CRs. The authors highlight the concept of moving the robot to create a denser map of the environment, but require external motion capture for localization. Our method takes a step towards the realization of the vision of [1] by introducing an algorithm for localizing a CR using sparse onboard ToF measurements without the need for external motion capture, marking the first demonstration of this capability on a soft continuum robot. We additionally move beyond the rigid-body assumption of [18] to consider the challenges of localizing a deformable continuum robot.

III. ESTIMATION FRAMEWORK

The estimation framework used in this work is based on the continuous-time factor-graph method presented in [32, 33]. The approach presents a factor-graph-based method to state estimation for CRs through a maximum a posteriori (MAP) optimization scheme. Prior factors are introduced that model a CR, with measurement factors being added for each sensor type. Factors are constructed using a weighted squared-error cost function. The total cost is minimized using a Gauss-Newton solver with uncertainties estimated via a Laplace approximation.

We use the same sliding-window formulation and prior factors, expanding upon the system with the introduction of new measurement factors described in the following sections. Our state $\mathbf{x}(s, t)$ consists of the continuous set of robot poses, $\mathbf{T}_{ib}(s, t) \in SE(3)$, body-centric generalized strains, $\boldsymbol{\epsilon}(s, t) \in \mathbb{R}^6$, and body-centric generalized velocities, $\boldsymbol{\varpi}(s, t) \in \mathbb{R}^6$ at time t and arc length s . $\mathbf{T}_{ib}(s, t)$ represents an $SE(3)$ transformation matrix that transforms points from a robot body frame, b , to an inertial frame, i . During optimization, the state is estimated at discrete times, t_k , and arc lengths, s_n , with the full continuous representation being available through the interpolation scheme described in [32].

A. Time-of-Flight Sensor

To include the ToF measurements into the factor-graph estimation framework, we follow the work of [6] and in-

TABLE I
TOF DISTANCE-BASED STANDARD DEVIATION MODEL.

Distance measurement [m]	Standard deviation [m]
$d < 0.025$	Measurement invalid
$0.025 \leq d < 0.6$	$0.014d \rightarrow 0.012d$
$0.6 \leq d < 1.2$	$0.012d \rightarrow 0.006d$
$1.2 \leq d$	$0.006d$

$A \rightarrow B$ indicates a linear interpolation from value A to value B over the specified distance range.

introduce a point-to-plane measurement factor with a Cauchy loss. Specifically, given a set of measured lidar distances $d_j \in \mathbb{R}$, $j = 0, 1, \dots, M$ from a sensor on the robot at time t and arc length s , we define a point-to-plane error as

$$e_j(\mathbf{x}) = \alpha \mathbf{n}_j^T \mathbf{D}(\mathbf{p}_j - \mathbf{T}_{ib}(t, s) \mathbf{q}_j^b) \quad (1)$$

where $\mathbf{p}_j$ is a matched point from a local map, $\mathbf{n}_j$ is the estimated normal at point $\mathbf{p}_j$, $\mathbf{D} \in \mathbb{R}^{3 \times 4}$ is a projection matrix that removes the homogeneous component, and $\alpha = (\sigma_2 - \sigma_3)/\sigma_1$ is a heuristic weight to favor planar neighborhoods. σ_i is the square root of the i th eigenvalue of the covariance obtained through principal component analysis of the neighborhood. See [11] for more details. $\mathbf{q}_j^b$ is a query point corresponding to measurement j and is assumed to be corrupted by isotropic Gaussian noise,

$$\mathbf{q}_j^b = \bar{\mathbf{q}}_j^b + \delta \mathbf{q}_j^b \quad \delta \mathbf{q}_j^b \sim \mathcal{N}(\mathbf{0}, \mathbf{D}^T \mathbf{R}_{\text{ToF}} \mathbf{D}). \quad (2)$$

Here, $\mathbf{R}_{\text{ToF}}$ is tuned experimentally starting with the relationships described in Table I with some additional noise to account for unmodeled process errors.

As ToF sensors are subject to disturbances that can lead to outlier measurements, we use a robust cost function to mitigate their effect on the estimation. As the estimation framework employed in this work uses a Gauss-Newton scheme designed for squared error factors, we implement a Cauchy loss via an iteratively reweighted least squares (IRLS) scheme [3, 17]. The final cost function for a single point-to-plane measurement is given by

$$J_j(\mathbf{x}) = \frac{1}{2} \mathbf{e}_j(\mathbf{x})^T \mathbf{Y}_j(\mathbf{x})^{-1} \mathbf{e}_j(\mathbf{x}), \quad \mathbf{Y}_j(\mathbf{x})^{-1} = \frac{\mathbf{R}_{\text{ToF}}^{-1}}{1 + \mathbf{e}_j(\mathbf{x})^T \mathbf{R}_{\text{ToF}}^{-1} \mathbf{e}_j(\mathbf{x})}. \quad (3)$$

As with other factors in the framework, we must compute the linearized error for the Gauss-Newton update step. This is given by

$$\mathbf{e}_j(\mathbf{x}) = \bar{\mathbf{e}}_j(\mathbf{x}) + \mathbf{G}_j|_{\mathbf{x}(s,t)} \delta \mathbf{t}_{bi}, \quad \mathbf{G}_j|_{\mathbf{x}(s,t)} = \alpha \mathbf{n}_j^T \mathbf{D}(\bar{\mathbf{T}}_{ib}(t, s) \mathbf{q}_j^b)^\odot \bar{\mathbf{T}}_{ib}(s, t). \quad (4)$$

This linearization is performed in an $SE(3)$ -sensitive way according to the methods described in [3]. $\delta \mathbf{t}_{bi} \in \mathbb{R}^6$ is the perturbation vector for the robot pose at time t and arc length s , $\bar{\mathbf{T}}_{ib}(s, t)$ is the adjoint matrix of the robot pose,

$\text{Ad}(\mathbf{T}_{ib}(s, t))$, and $(\cdot)^\odot$ is the operator defined as

$$\begin{bmatrix} \mathbf{a} \\ b \end{bmatrix}^\odot = \begin{bmatrix} b \mathbf{I}_{3 \times 3} & -\mathbf{a}^\wedge \\ \mathbf{0}_{1 \times 3} & 0 \end{bmatrix}, \quad (5)$$

$$\mathbf{a}^\wedge = \begin{bmatrix} 0 & -a_3 & a_2 \\ a_3 & 0 & -a_1 \\ -a_2 & a_1 & 0 \end{bmatrix}. \quad (6)$$

B. Gyroscope Sensor

In addition to the ToF sensors, gyroscope measurements are also included in the factor graph to aid in localization. The gyroscope measurements are modeled as in [32]. As we are working on small time scales, gyroscope biases are approximated to be constant and are determined by averaging an initial stationary bias.

C. Strain Sensor (Simulation Only)

The strain data used in this work has a bending angle and curvature representation that is consistent with many FBG sensing systems, enabling future hardware integration in more complex settings. Bending angle θ_j and curvature κ_j are mapped to a measured strain value $\tilde{\epsilon}_j$ via

$$\tilde{\epsilon}_j = \mathcal{T}(\theta_j) [1 \ 0 \ 0 \ 0 \ 0 \ \kappa_j]^T, \quad (7)$$

$$\mathcal{T}(\theta_j) = \text{Ad} \left(\begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos(\theta_j) & -\sin(\theta_j) & 0 \\ 0 & \sin(\theta_j) & \cos(\theta_j) & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \right). \quad (8)$$

An associated error factor is used via the cost function

$$J_j = \frac{1}{2} (\epsilon_j - \tilde{\epsilon}_j)^T \mathbf{R}_{\text{strain}}^{-1} (\epsilon_j - \tilde{\epsilon}_j), \quad (9)$$

where $\mathbf{R}_{\text{strain}}$ is the strain measurement covariance and is modeled to be isotropic.

D. Prior Map

Each localization instance begins with a prior map of the environment represented as a point cloud with precomputed normals and planarity weights. For simulated environments, this map is generated via Poisson sampling of a CAD-generated scene. For real-world scenes, a point cloud is generated via scanning the scene with a CreaLity Raptor 3D scanner [31]. Reflective markers are distributed in the scene for coordinate alignment to a motion capture system, and the near-infrared scanning mode is adopted, achieving a 0.1 mm accuracy with a 0.1 mm to 2 mm resolution. The map is stored in a hashed voxel map to enable fast nearest neighbor queries [6], though other representations (e.g., KD trees, PH-trees) are compatible with the work presented here. By varying which prior map the estimator has access to while keeping the sensor data the same, we can create scenarios where differences between the prior map and the scene are introduced. This is how deviations to the environment are created in both simulation and real-world experiments.

E. Scene Reconstruction

After estimation is complete, the robot's trajectory and ToF measurements are used to reconstruct a point cloud of the environment. Each ToF measurement can be projected into the inertial frame using the estimated robot pose at the time of the measurement using the uncertainty-aware projection,

$$\bar{p}_j = \bar{T}_{ib}(t, s) \bar{q}_j^b, \quad (10)$$

$$\delta p_j = \begin{bmatrix} -\bar{p}_j^\odot \bar{T}_{ib} & \bar{T}_{ib} \end{bmatrix} \begin{bmatrix} \delta t_{bi} \\ \delta q_j^b \end{bmatrix} \sim \mathcal{N}(\mathbf{0}, D^T \Sigma_{p_j} D),$$

$$\Sigma_{p_j} = R_{\text{ToF}} + D \bar{p}_j^\odot \bar{T}_{ib} \Sigma_b \bar{T}_{ib}^T \bar{p}_j^\odot D^T. \quad (11)$$

Here, Σ_b is the body-frame pose covariance as estimated during the optimization process. The expression in (11) makes use of the fact that R_{ToF} is isotropic.

IV. SIMULATION

For generating simulated data, we use the MuJoCo physics engine [34] with a number of custom additions to simulate a CR operating in a confined environment with ToF, gyroscope, and strain sensors. The simulated scene is designed to emulate the real-world application of jet engine inspection. The scene dimensions, robot specs, and anomaly choices are meant to reflect this.

A. Robot Details

The simulated robot is modeled in MuJoCo as a series of rigid links connected by spherical joints. This approximation follows the pseudo-rigid body method commonly used in CR literature [23]. Provided the links are short enough, this robot will approximate the behavior of a true CR, while allowing it to take advantage of MuJoCo's extensive contact and dynamics modeling capabilities. The robot has three tendon-actuated modules with a combined length of 1.5 m, with each MuJoCo link being 1 cm long. Every 30 cm is a sensor disc containing three radially mounted ToF sensors and one gyroscope sensor. The discs are offset by one another by an axial rotation of 60°. The tip of the robot contains an additional ToF sensor pointing forward. A strain sensor is simulated down the center of the robot, providing bending strain measurements every 3 cm. See section IV-C for more information.

B. Environment Setup

Motivated by the task of aircraft inspection [10, 12, 37], the environment used in this work comes from a model jet engine [9]. The model is scaled to have an outer diameter of 3.1 m (approximately the size of a commercial airliner engine) and is cropped to contain only the front compressor section. The mesh is down-sampled to adhere to MuJoCo's 200,000 face limit on imported objects. A small section of the mesh where robot-environment collisions are expected is decomposed into a series of convex hulls using CoACD [38] to improve collision performance. Gravity, friction, and collisions are all enabled in the simulation. Fig. 2 shows the robot and environment used in the simulation.




Fig. 2. Simulated jet engine environment with CR in the MuJoCo physics engine. The robot is 1.5 m long and contains distributed ToF, IMU, and strain sensors. The environment mesh is displayed in blue while the collision mesh is shown in green. Time-of-flight sensor rays are visualized in yellow.

Two types of map differences are added to the environment: unmodeled objects and missing modeled features. Unmodeled objects are simulated by adding additional meshes into the environment to be sensed and collided with by the robot. Missing modeled features are simulated by selectively removing parts of the environment mesh. In total, 10 environments are created: three with unmodeled objects (scenes S1–S3), three with missing modeled features (scenes S4–S6), three with both types of anomalies (scenes S7–S9), and one with no changes at all (scene S0). The anomalies range in size and are placed in locations that are reachable by the robot.

C. Sensor Simulation

1) *Time-of-flight Sensor*: The ToF sensors are designed to simulate the VL53L5CX sensors from STMicro-Electronics used on the real robot in this work. It is simulated by a grid of 8×8 MuJoCo rangefinder sensors with a shared origin and field of view of 45° both horizontally and vertically. Each rangefinder has a maximum range of 4 m. Noise is added to each range measurement according to the model described in Table I. Each ToF sensor generates data at a rate limited to at most 15 Hz.

2) *Gyro Sensor*: The gyro sensors are modeled using the Mujoco gyroscope sensor. Noise is added to the gyroscope measurements using a zero-mean Gaussian with a standard deviation of 0.01 rad/s. The sensor data is sampled at a rate of 100 Hz.



Fig. 3. Experimental setup including the extensible continuum robot, the OptiTrack camera system, and a close-up of one of the sensorized rings carrying the markers, ToF sensors, and IMUs. The scene depicted contains all four objects present (scene R4).

3) *Strain Sensor*: Strain sensing is simulated through a custom MuJoCo sensor plug-in by computing the bending angle and curvature at specified points along the robot. To map from the pseudo-rigid body representation of the robot to a continuous strain profile a constant-curvature interpolation scheme is employed. Specifically, each link's midpoint is treated as the endpoint of a constant-curvature arc whose ends are tangent to the rigid link at the midpoint. This piecewise constant-curvature approximation enables strain estimation at any point along the robot. Zero-mean Gaussian noise with standard deviations of 0.01 m^{-1} and 0.015 rad is added to the curvature and bending-angle measurements, respectively. Strain data is sampled at a rate of 20 Hz.

V. EXPERIMENTAL SETUP

The experimental setup, shown in Fig. 3, comprises a cube-shaped environment ($70 \text{ cm} \times 70 \text{ cm} \times 60 \text{ cm}$ W $\times$ D $\times$ H), four OptiTrack cameras, a set of objects and a soft, continuum

robot arm. The arm consists of three modules, each having three flexures printed out of Thermoplastic Polyurethane (TPU) 82A (Recreus) and a sensorized ring containing three ToF sensors and a gyroscope unit. Both the flexures and the ToFs for each module were oriented at 120° from one another, and each module was offset by 60° with respect to the one above it. At rest, the robot measures in at 53 cm in length, with each module being approximately the same length. The experimental scene is designed to stress-test the limits of the proposed algorithm, with progressive anomalies increasing in size until localization failure is achieved.

A. Sensor Characteristics

The ToF sensors used in this work are the VL53L5CX sensors from STMicro-Electronics. Each sensor has a horizontal and vertical field of view of 45° and produces an 8×8 grid of ToF measurements. The sensor noise characteristics of the VL53L5CX sensor were experimentally determined in [8], where it was found that the sensor noise can be modeled as a series of Gaussian distributions dependent on the measured range. Specifically, the measurement noise standard deviation for a single measurement is given in Table I. Each ring also incorporates an ISM330BX IMU (STMicro-Electronics) setup with an accelerometer full-scale of $\pm 4 \text{ g}$ and a gyroscope full-scale range of $\pm 500 \text{ dps}$. For this work, only the gyroscope measurements are used in the estimation framework. The IMU is sampled at 120 Hz while the ToF sensors are sampled at 15 Hz.

B. Data Collection and Environments

A total of five environments were created for experimental validation, each with one more object added to the scene than the previous. The objects added in order are: a mustard bottle (scene R1), a cup (scene R2), a cracker box (scene R3), and a wet floor sign (scene R4), with the final environment being a scene with no objects (scene R0). The objects range in size and shape to provide a variety of conditions for testing. Each environment is scanned using the 3D scanner described in section III-D to create a prior map for localization. A total of 23 trajectories are collected across the five environments, with all having at least four trajectories. During each trajectory, four tendons actuated the robot: one in the central module and three in the distal one. Due to the compliance of the flexures, the arm displays significant compression when being actuated, rendering any Kirchhoff rod approximations infeasible.

VI. RESULTS AND DISCUSSION

For all results shown (real-world and simulation), the estimator is initially tuned using one hold-out trial consisting of a single trajectory operated in a fully known environment (R4 for real-world and S0 for simulation). Tuning is done by sweeping through reasonable values for the process and measurement noise parameters and selecting the set that minimizes the average localization error at the marker locations. The same tuned parameters are then used for all other trials.

TABLE II
LOCALIZATION ERROR IN SIMULATED ENVIRONMENT USING SCENE S0 AS
A PRIOR MAP. VALUES SHOWN AS MEAN $\pm$ STANDARD DEVIATION
ACROSS ALL TRIALS AND RINGS.

Scene	Translation Error (cm)		Rotation Error (deg)	
	MAE	RMSE	MAE	RMSE
S0	0.78 ± 0.48	0.99 ± 0.65	0.96 ± 0.48	0.98 ± 0.48
S1	0.78 ± 0.48	0.99 ± 0.65	0.96 ± 0.48	0.98 ± 0.48
S2	0.78 ± 0.47	1.00 ± 0.64	0.96 ± 0.49	0.98 ± 0.48
S3	0.78 ± 0.48	0.99 ± 0.65	0.96 ± 0.48	0.98 ± 0.48
S4	0.61 ± 0.34	0.74 ± 0.38	0.96 ± 0.46	0.98 ± 0.45
S5	0.73 ± 0.43	0.94 ± 0.58	0.96 ± 0.48	0.98 ± 0.48
S6	0.77 ± 0.47	0.98 ± 0.64	0.96 ± 0.48	0.98 ± 0.48
S7	0.61 ± 0.34	0.74 ± 0.38	0.96 ± 0.46	0.98 ± 0.45
S8	0.76 ± 0.44	0.96 ± 0.60	0.96 ± 0.49	0.98 ± 0.48
S9	0.77 ± 0.47	0.98 ± 0.64	0.96 ± 0.48	0.98 ± 0.48

A. Simulation

To isolate the effect of prior map accuracy on localization performance, we conducted a series of simulation experiments where the robot actuation was kept identical while varying the scenes. Three robot trajectories are executed in each of the 10 scenes (S0–S9). For each of the 30 trials, the prior map used was that of scene S0. The results are summarized in Table II. The simulation results show no noticeable degradation in localization performance when the prior map deviates from the true environment due to the addition or removal of geometries from the map. This is made possible through the use of the robust cost function employed with the ToF factors in this work. The accuracy shown in the simulation deviates greatly from the real-world results, primarily due to the role of the strain sensor data in the estimate, as well as the idealized nature of the simulated environment and comparatively small map deviations, among other sources of unmodeled error. The strain sensor alone, in the absence of time-of-flight, does not provide enough information for accurate localization, resulting in ‘spatial drift’ in the estimates. In this case, the error in the estimate increases from the base towards the tip, resulting in a large dip in tip accuracy compared to the fused time-of-flight and strain sensor setup. Likewise, time-of-flight without strain measurements results in estimates that frequently converge to local minima, due in large part to the radial symmetry of the environment.

Side-by-side comparisons of the estimated robot shape against the simulated robot for the S8 scene are shown in Fig. 4. Fig. 5 shows a side-by-side of the true scene, prior map, and reconstructed scene via ToF measurements and the method described in Section III-E. The scene is constructed using all estimates collected on scene S8. In the simulated environment, the small anomalies present are not large enough to significantly affect localization and scene reconstruction accuracy.

B. Real World Experiments

Localization performance is evaluated for every combination of prior map and real environment (that varies by one




Fig. 4. Side-by-side comparison of the simulated environment (left) and the estimated robot (right). The estimated robot shape and trajectory closely match the ground truth, with an average position error at the marker locations of less than 1 cm and rotational error of less than 1°. Anomalies in the environment are marked—they appear to have little effect on localization performance in simulation.

object) as described in Section V-B. The results are summarized in Table III. For trials where the prior map matches the real environment, the estimator performs accurately and consistently, achieving average position errors between 1.2 cm to 2.5 cm and rotational ones between 3.7° to 10.3°.

When we introduce deviation between the true scenes and the prior, localization performance typically degrades. When an object is added to the environment, the localization error increases by an average of 1.7 cm and 1.0° across all trials. When an object is removed from the environment, the localization error increases by an average of 0.5 cm and 1.5°. The largest degradation in performance occurs when the large wet-floor sign is added to the environment, causing the average position error to increase to 5.1 cm and rotational error to 6.5° when added and to 3.6 cm in position and 14.6° in rotation when it is removed.

The trial with the median performance (of all 61 combinations tested) is highlighted in Fig. 6, showing a side-by-side




Fig. 5. Simulated scene S8 (left), prior map S0 (middle), and the reconstructed scene (right) via ToF measurements. A depth-based color gradient is applied to the point clouds for visual clarity. The small objects added and removed between the two scenes appear to have little effect on localization performance in simulation.

TABLE III

LOCALIZATION ERROR GROUPED BY EXPERIMENT CONDITION. VALUES SHOWN AS MEAN $\pm$ STANDARD DEVIATION ACROSS ALL TRIALS AND MOTION CAPTURE RINGS.

Map – Scene	Translation Error (cm)		Rotation Error (deg)	
	MAE	RMSE	MAE	RMSE
Localization within known environment (no anomalies)				
R0 – R0	1.8 ± 0.6	1.9 ± 0.7	5.2 ± 1.4	5.8 ± 1.8
R1 – R1	1.9 ± 0.9	2.3 ± 1.3	6.3 ± 4.3	7.8 ± 6.0
R2 – R2	2.0 ± 1.0	2.4 ± 1.3	6.3 ± 3.9	7.9 ± 5.5
R3 – R3	1.7 ± 0.8	1.9 ± 1.0	5.9 ± 2.6	6.8 ± 3.7
R4 – R4	1.9 ± 0.4	2.0 ± 0.4	8.4 ± 5.0	9.1 ± 5.1
Localization in environments with an additional object				
R0 – R1	3.2 ± 2.0	3.5 ± 2.1	8.6 ± 6.6	9.7 ± 7.3
R1 – R2	2.6 ± 1.8	3.0 ± 2.3	7.3 ± 5.3	9.0 ± 7.3
R2 – R3	3.2 ± 2.2	3.5 ± 2.2	7.0 ± 3.2	8.0 ± 4.1
R3 – R4	5.1 ± 2.1	5.3 ± 2.1	6.5 ± 4.3	7.6 ± 5.8
Localization in environments with a removed object				
R1 – R0	1.7 ± 0.6	1.9 ± 0.7	4.3 ± 1.5	4.9 ± 2.1
R2 – R1	2.0 ± 1.0	2.4 ± 1.4	6.4 ± 4.4	7.8 ± 6.1
R3 – R2	2.0 ± 1.0	2.4 ± 1.3	6.4 ± 3.8	7.9 ± 5.4
R4 – R3	3.6 ± 2.6	4.0 ± 2.9	14.6 ± 4.0	16.8 ± 5.3

comparison of the real robot and the estimated robot shape at several timesteps. The estimator clearly tracks the robot shape well, with this particular trial achieving an average position error of 1.7cm and rotational error of 6.7° at the marker locations.

As with most continuum robot applications, the tip ring tends to have the worst localization performance due to error accumulation along the length of the robot. Average localization errors for each ring across all trials are summarized in Table IV. While the ToF readings directly constrain the tip position along some axes, the distance from the strong prior information at the base of the robot leads to the tip disc being most susceptible to local minima.

TABLE IV
LOCALIZATION ERROR OF INDIVIDUAL MOTION CAPTURE RINGS. RESULTS POOLED ACROSS ALL EXPERIMENT CONDITIONS.

Ring	Translation MAE (cm)	Rotation MAE (deg)
Top	1.7 ± 1.3	4.8 ± 4.1
Middle	2.3 ± 1.2	7.0 ± 3.4
Bottom (Tip)	3.6 ± 1.9	9.9 ± 4.9



A side-by-side comparison of the real environment, prior map, and reconstructed scene via ToF measurements for the R4 – R4, R3 – R4, and R3 – R3 trials is shown in Fig. 7. In the cases where the prior map is accurate to the scene, reconstruction closely matches the real environment (per-point RMSE of 0.8 cm for R4 – R4, 1.2 cm for R3 – R3). When large anomalies degrade the localization performance, the reconstructed scene also deviates significantly from the real environment (per-point RMSE of 1.4 cm). This indicates that improving localization performance in the presence of anomalies will also improve scene reconstruction accuracy.

C. Limitations and Future Work

1) *Sensor Degeneracies and Local Minima*: As the ToF sensors acquire sparse measurements of the environment, even in feature-rich scenes, they are more susceptible to causing




Fig. 7. Real scenes (left), prior maps (middle), and the reconstructed scenes (right) via ToF measurements for the R3 – R3 (top), R3 – R4 (middle), and R4 – R4 scenarios. A depth-based color gradient is applied to the point clouds for visual clarity. Objects in the scene with high reflectance or low incidence angle to the ToF sensors can lead to noisy measurements that affect reconstruction quality. This can be seen most clearly surrounding the cup on the left side of both scenes.

the estimator to get stuck in local minima. Without a method for detecting and escaping these minima, the estimator can produce large localization errors in certain scenarios. Example failure cases are shown in Fig. 8. A poor initialization, significant departure of the real scene from the map, and sudden movements can all lead to the estimator more easily converging to local minimum solutions. This is not too surprising given the information being generated by the real-world robot. The experiment setup used in this work is particularly challenging with respect to degenerate scenarios due to the robot being extensible, the sensors having a limited field of view, and the scene having many symmetries. As demonstrated in the simulation result, the addition of more diverse sensor modalities (e.g., strain) helps alleviate this issue. The strain sensor information provides a restoring force effect to solutions that would otherwise have jumped to a local minimum. The ISM330BX used as a gyroscope in this work also provides accelerometer measurements that could be incorporated to aid in this regard, though this is left to future work.

2) *Estimator Smoothing and Tradeoffs*: The estimator framework used in this work introduces both temporal and spatial smoothing of the estimate. This effect is a double-edged sword. On one hand, it helps regularize the estimate and avoid overfitting to noisy measurements. On the other hand, it can prevent the estimator from accurately capturing rapid changes in the robot’s shape over time. The local minima observed in the failure cases appear to be a result of both a highly deviated scene from the prior map and a sudden change in robot shape.



Fig. 8. Example side-by-side comparison of the real environment (left) and the estimated robot (right) for failure cases in each of the three worst-performing trials (highest average localization error). In each case, the estimated robot shape and trajectory deviate significantly from the ground truth, getting stuck in local minima. Each of these scenarios features a map that is particularly challenging due to its deviation from the real system, where the large wet-floor sign has been removed.

3) *Prior Map Dependency*: As shown clearly in the results, the localization accuracy of this method is currently tied to the accuracy of the prior map. The larger this deviation, the worse the localization performance. While this is expected, it does limit the applicability of this method to scenarios where a reasonably accurate prior map can be obtained. Future work could explore methods for relaxing this requirement, incorporating live map updates, or otherwise improving robustness to map inaccuracies.

4) *Application: Anomaly Detection*: One possible application of this work is anomaly detection in support of inspection tasks. To detect anomalies, each point in the reconstructed point cloud is compared against the prior map. A point is classified as an anomaly if its squared Mahalanobis distance to the nearest neighbor in the prior map exceeds a threshold τ , i.e.

$$(\bar{\mathbf{p}}_j - \mathbf{p}_{nn})^T (\Sigma_{p_j} + \Sigma_{nn})^{-1} (\bar{\mathbf{p}}_j - \mathbf{p}_{nn}) > \tau, \quad (12)$$

where $\mathbf{p}_{nn}$ is the nearest neighbor point in the prior map with prior covariance of Σ_{nn} . Fig. 9 highlights some results of this detector using the data from this work. With the simulated scenes performing slightly better with this rejection technique, additional work is required to improve the localization to the




Fig. 9. Real scene (top) with detected anomalies shown over prior map (bottom) highlighted in red. The detection is more successful in the simulated environment, where localization tends to be more accurate. Real-world experiments show false-positive detections near scene objects that corrode sensor readings. Further improvements to noise reduction and localization accuracy are required to make this a viable anomaly detection method.

point of being able to reliably detect the changes in the real-world experiments.

VII. CONCLUSION

In this work, we show that full-body localization of continuum robots is achievable using sparse, low-resolution ToF sensors distributed along the robot body. Our proposed state estimation framework achieves an average localization accuracy of 2.5 cm on an extensible continuum robot in 3D space, despite frequent sensing degeneracies and moderate deviations between the prior map and the true scene. These results demonstrate that compact, body-mounted depth sensors can support practical continuum robot localization without external tracking, while highlighting remaining challenges related to prior map dependence and local minima that motivate future work.

ACKNOWLEDGMENTS

We acknowledge the support of the Natural Sciences and Engineering Research Council of Canada (NSERC) [funding reference number RGPIN-2025-04343]. This research was also supported by the Canada Foundation for Innovation (CFI) through the John R. Evans Leaders Fund (JELF) [project number 40110]. This work was also supported by SESTOSENSE (HORIZON EUROPE Research and Innovation Actions) under GA number 101070310 and by the Engineering and Physical Sciences Research Council (EPSRC) Grant EP/V000748/1. We would like to thank Chengnan (Jimmy) Shentu and Nicholas Baldassini for their guidance in constructing environments and continuum robots in MuJoCo.

REFERENCES

[1] Colette Abah, Andrew L. Orekhov, Garrison L. H. Johnston, and Nabil Simaan. A multi-modal sensor array for human–robot interaction and confined spaces exploration using continuum robots. *IEEE Sensors Journal*, 22(4): 3585–3594, 2022. doi: 10.1109/JSEN.2021.3140002.

[2] GE Aerospace. Blade inspection toolkit. URL <https://www.geaerospace.com/commercial/services/engine-maintenance-technologies/blade-inspection-tool>.

[3] T D Barfoot. *State Estimation for Robotics*. Cambridge University Press, 2017. ISBN 9781107159396. doi: 10.1017/9781316671528.

[4] Justin A. Borgstadt, Michael R. Zinn, and Nicola J. Ferrier. Multi-modal localization algorithm for catheter interventions. In *2015 IEEE International Conference on Robotics and Automation (ICRA)*, pages 5350–5357, 2015. doi: 10.1109/ICRA.2015.7139946.

[5] Jessica Burgner-Kahrs, D. Caleb Rucker, and Howie Choset. Continuum robots for medical applications: A survey. *IEEE Transactions on Robotics*, 31(6):1261–1280, 2015. doi: 10.1109/TRO.2015.2489500.

[6] Keenan Burnett, Angela P. Schoellig, and Timothy D. Barfoot. Continuous-time radar-inertial and lidar-inertial odometry using a gaussian process motion prior. *IEEE Transactions on Robotics*, 41:1059–1076, 2025. doi: 10.1109/TRO.2024.3521856.

[7] Giammarco Caroleo, Alessandro Albin, and Perla Maiolino. Soft robot localization using distributed miniaturized time-of-flight sensors. In *2025 IEEE 8th International Conference on Soft Robotics (RoboSoft)*, pages 1–6. doi: 10.1109/RoboSoft63089.2025.11020858. ISSN: 2769-4534.

[8] Giammarco Caroleo, Alessandro Albin, Daniele De Martini, Timothy D. Barfoot, and Perla Maiolino. Tiny lidars for manipulator self-awareness: Sensor characterization and initial localization experiments. In *2025 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 12940–12946, 2025. doi: 10.1109/IROS60139.2025.11246575.

[9] CATIAV5FTW. “3d printable jet engine” [3d model]. <https://www.printables.com/model/572421-3d-printable-jet-engine>, 2025. Adapted by the authors. Original work licensed under Creative Commons Attribution-NonCommercial 4.0 International (CC BY-NC 4.0).

[10] Jung-Che Chang, Hengtai Dai, Xi Wang, Dragos Axinte, and Xin Dong. Development of a continuum robot with inflatable stiffness-adjustable elements for in-situ repair of aeroengines. *Robotics and Computer-Integrated Manufacturing*, 95:103018, 2025. ISSN 0736-5845. doi: <https://doi.org/10.1016/j.rcim.2025.103018>.

[11] J. Demantké, C. Mallet, N. David, and B. Vallet. Dimensionality based scale selection in 3d lidar point clouds. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, XXXVIII-5/W12:97–102, 2011.

[12] X. Dong, D. Axinte, D. Palmer, S. Cobos, M. Raffles, A. Rabani, and J. Kell. Development of a slender continuum robotic system for on-wing inspection/repair of gas turbine engines. *Robotics and Computer-Integrated Manufacturing*, 44:218–229, 2017. ISSN 0736-5845. doi: <https://doi.org/10.1016/j.rcim.2016.09.004>.

- [13] Pierre E. Dupont, Nabil Simaan, Howie Choset, and Caleb Rucker. Continuum robots for medical interventions. *Proceedings of the IEEE*, 110(7):847–870, 2022. doi: 10.1109/JPROC.2022.3141338.
- [14] James M. Ferguson, D. Caleb Rucker, and Robert J. Webster. Unified shape and external load state estimation for continuum robots. *IEEE Transactions on Robotics*, 40:1813–1827, 2024. doi: 10.1109/TRO.2024.3360950.
- [15] Alfred M. Franz, Tamás Haidegger, Wolfgang Birkfellner, Kevin Cleary, Terry M. Peters, and Lena Maier-Hein. Electromagnetic tracking in medicine—a review of technology, validation, and applications. *IEEE Transactions on Medical Imaging*, 33(8):1702–1725, 2014. doi: 10.1109/TMI.2014.2321777.
- [16] Cedric Girerd, Andrey V. Kudryavtsev, Patrick Rougeot, Pierre Renaud, Kanty Rabenorosoa, and Brahim Tamadazte. Slam-based follow-the-leader deployment of concentric tube robots. *IEEE Robotics and Automation Letters*, 5(2):548–555, 2020. doi: 10.1109/LRA.2019.2963821.
- [17] Paul W. Holland and Roy E. Welsch. Robust regression using iteratively reweighted least-squares. *Communications in Statistics - Theory and Methods*, 6(9):813–827, 1977. doi: 10.1080/03610927708827533.
- [18] Kengo Iwao, Hikaru Arita, and Kenji Tahara. State estimation and environment recognition for articulated structures via proximity sensors distributed over the whole body. *IEEE Robotics and Automation Letters*, 10(3):3030–3037, 2025. doi: 10.1109/LRA.2025.3539117.
- [19] Beobkyyoon Kim, Junhyoung Ha, Frank C. Park, and Pierre E. Dupont. Optimizing curvature sensor placement for fast, accurate shape sensing of continuum robots. In *2014 IEEE International Conference on Robotics and Automation (ICRA)*, pages 5374–5379, 2014. doi: 10.1109/ICRA.2014.6907649.
- [20] Aditya Brij Koolwal, Federico Barbagli, Christopher Carlson, and David Liang. An ultrasound-based localization algorithm for catheter ablation guidance in the left atrium. *The International Journal of Robotics Research*, 29(6):643–665, 2010. doi: 10.1177/0278364909105332.
- [21] Sven Lilge, Timothy D. Barfoot, and Jessica Burgner-Kahrs. Continuum robot state estimation using gaussian process regression on se(3). *The International Journal of Robotics Research*, 41(13-14):1099–1120, 2022. doi: 10.1177/02783649221128843.
- [22] Arthur W. Mahoney, Trevor L. Bruns, Philip J. Swaney, and Robert J. Webster. On the inseparable nature of sensor selection, sensor placement, and state estimation for continuum robots or “where to put your sensors and how to use them”. In *2016 IEEE International Conference on Robotics and Automation (ICRA)*, pages 4472–4478. doi: 10.1109/ICRA.2016.7487646.
- [23] Priyanka Rao, Quentin Peyron, Sven Lilge, and Jessica Burgner-Kahrs. How to model Tendon-Driven continuum robots and benchmark modelling performance. *Front Robot AI*, 7:630245, February 2021.
- [24] Austin Reiter, Roger E. Goldman, Andrea Bajo, Konstantinos Iliopoulos, Nabil Simaan, and Peter K. Allen. A learning algorithm for visual pose estimation of continuum robots. In *2011 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 2390–2396, 2011. doi: 10.1109/IROS.2011.6094947.
- [25] William S. Rone and Pinhas Ben-Tzvi. Multi-segment continuum robot shape estimation using passive cable displacement. In *2013 IEEE International Symposium on Robot and Sensors Environments (ROSE)*, pages 37–42, 2013. doi: 10.1109/ROSE.2013.6698415.
- [26] Matteo Russo, Seyed Mohammad Hadi Sadati, Xin Dong, Abdelkhalick Mohammad, Ian D. Walker, Christos Bergeles, Kai Xu, and Dragos A. Axinte. Continuum robots: An overview. *Advanced Intelligent Systems*, 5(5):2200367, 2023. doi: <https://doi.org/10.1002/aisy.202200367>.
- [27] Chengnan Shentu, Enxu Li, Chaojun Chen, Puspita T Dewi, David B Lindell, and Jessica Burgner-Kahrs. Moss: Monocular shape sensing for continuum robots. *IEEE Robotics and Automation Letters*, 9(2):1524–1531, 2023.
- [28] Chaoyang Shi, Xiongbiao Luo, Peng Qi, Tianliang Li, Shuang Song, Zoran Najdovski, Toshio Fukuda, and Hongliang Ren. Shape sensing techniques for continuum robots in minimally invasive surgery: A survey. *IEEE Transactions on Biomedical Engineering*, 64(8):1665–1678, 2017. doi: 10.1109/TBME.2016.2622361.
- [29] Peter Jan Sincak, Erik Prada, Ľubica Miková, Roman Mykhailishyn, Martin Varga, Tomas Merva, and Ivan Virgala. Sensing of continuum robots: A review. *Sensors*, 24(4), 2024. ISSN 1424-8220. doi: 10.3390/s24041311.
- [30] Christian Sorensen, Phillip Hyatt, Matthew Ricks, Seth Nielsen, and Marc D. Killpack. Soft robot configuration estimation and control using simultaneous localization and mapping. In *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 616–623, 2021. doi: 10.1109/IROS51168.2021.9635896.
- [31] Crealitiy Store. Cr-scan raptor 3d scanner, 2026. URL <https://store.crealitiy.com/uk/products/cr-scan-raptor-3d-scanner>.
- [32] Spencer Teetaert, Sven Lilge, Jessica Burgner-Kahrs, and Timothy D. Barfoot. A stochastic framework for continuous-time state estimation of continuum robots. 2025. *arXiv:2510.01381*.
- [33] Spencer Teetaert, Sven Lilge, Jessica Burgner-Kahrs, and Timothy D. Barfoot. A sliding-window filter for online continuous-time continuum robot state estimation. In *2026 IEEE 9th International Conference on Soft Robotics (RoboSoft)*, pages 239–246, 2026.
- [34] Emanuel Todorov, Tom Erez, and Yuval Tassa. Mujoco: A physics engine for model-based control. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 5026–5033. IEEE, 2012. doi: 10.1109/IROS.2012.6386109.
- [35] T. S. Vaquero, G. Daddi, R. Thakker, M. Paton, A. Jasour,

- M. P. Strub, R. M. Swan, R. Royce, M. Gildner, P. Tosi, M. Veismann, P. Gavrilo, E. Marteau, J. Bowkett, D. Loret de Mola Lemus, Y. Nakka, B. Hockman, A. Orekhov, T. D. Hasseler, C. Leake, B. Nuernberger, P. Proença, W. Reid, W. Talbot, N. Georgiev, T. Pailevanian, A. Archanian, E. Ambrose, J. Jasper, R. Etheredge, C. Roman, D. Levine, K. Otsu, S. Yearicks, H. Melikyan, R. R. Rieber, K. Carpenter, J. Nash, A. Jain, L. Shiraishi, M. Robinson, M. Travers, H. Choset, J. Burdick, A. Gardner, M. Cable, M. Ingham, and M. Ono. Eels: Autonomous snake-like robot with task and motion planning capabilities for ice world exploration. *Science Robotics*, 9(88):eadh8332, 2024. doi: 10.1126/scirobotics.adh8332.
- [36] Xiaomei Wang, Jing Dai, Hon-Sing Tong, Kui Wang, Ge Fang, Xiaochen Xie, Yun-Hui Liu, Kwok Wai Samuel Au, and Ka-Wai Kwok. Learning-based visual-strain fusion for eye-in-hand continuum robot pose estimation and control. *IEEE Transactions on Robotics*, 39(3):2448–2467, 2023. doi: 10.1109/TRO.2023.3240556.
- [37] Yaming Wang, Feng Ju, Yahui Yun, Jiafeng Yao, Yaoyao Wang, Hao Guo, and Bai Chen. An inspection continuum robot with tactile sensor based on electrical impedance tomography for exploration and navigation in unknown environment. *The Industrial Robot*, 47(1):121–130, 2020.
- [38] Xinyue Wei, Minghua Liu, Zhan Ling, and Hao Su. Approximate convex decomposition for 3d meshes with collision-aware concavity and tree search. *ACM Transactions on Graphics (TOG)*, 41(4):1–18, 2022.
- [39] Yu Yamauchi, Yuichi Ambe, Hikaru Nagano, Masashi Konyo, Yoshiaki Bando, Eisuke Ito, Solvi Arnold, Kimitoshi Yamazaki, Katsutoshi Itoyama, Takayuki Okatani, Hiroshi G Okuno, and Satoshi Tadokoro. Development of a continuum robot enhanced with distributed sensors for search and rescue. *ROBOMECH J.*, 9(1), December 2022.
- [40] Jingyu Zhang, Qin Fang, Pingyu Xiang, Danying Sun, Yanan Xue, Rui Jin, Ke Qiu, Rong Xiong, Yue Wang, and Haojian Lu. A survey on design, actuation, modeling, and control of continuum robot. *Cyborg and Bionic Systems*, 2022:9754697, 2022. doi: 10.34133/2022/9754697.