2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)
September 27 - October 1, 2021. Prague, Czech Republic

Soft Robot Configuration Estimation and Control Using Simultaneous Localization and Mapping

Christian Sorensen* Phillip Hyatt* Matthew Ricks* Seth Nielsen* Marc D. Killpack

2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) | 978-1-6654-1714-3/21/$31.00 ©2021 IEEE | DOI: 10.1109/IROS51168.2021.9635896

***Abstract*— In this paper we present a novel approach to accomplishing soft robot configuration estimation and control using RGB-D cameras and SLAM-based methods. By placing cameras on the unactuated sections of our large-scale (approximately 2 meters long) pneumatic soft robot, we can map an environment and then estimate the orientation of the robot links using landmark-based localization. Using the orientations of each camera we can solve for the joint configurations between them. We first show that this method works for a traditional rigid robot (Baxter) where we can compare against the ground truth encoder values. For Baxter, the median joint angle error was on the order of 1-2°. We then show that the SLAM-based method provides estimates for soft robot configuration that are within 1° when compared to our past methods of using a HTC Vive Tracker. While HTC Vive Trackers and commonly used motion capture systems require externally mounted sensors placed in the robot's environment, the SLAM-based estimation method presented here works in any visually feature-rich environment. Finally we show that this method of estimation is effective for closed-loop control of soft robots by controlling our large-scale soft robot through a series of joint configurations.**

I. Introduction

Despite the promise of soft robots being able to fundamentally change the way that robots interact with the world, this promise will not be fully realized without soft robots being able to estimate their configuration in an onboard, online manner. Without this kind of estimation, effective closed-loop control outside of laboratory environments is improbable. This soft robot estimation problem is multi-faceted and will likely not be solved by a single sensing technology or approach. However, in this paper we present a novel method for online configuration estimation by using cameras in a simultaneous localization and mapping (SLAM) framework.

Specifically, in this paper we use RGB-D cameras with an existing SLAM library (ORB-SLAM2 [1]). We use SLAM to estimate the orientations of two cameras fixed to the soft robot links relative to an initial starting configuration. Using assumptions about the starting configuration, as well as camera orientations supplied by SLAM, the configuration of the soft robot joints can be estimated.

The method presented in this paper has strengths in areas where many other soft robot sensing technologies are lacking. Specifically, most existing technologies do not localize relative to the real-world environment, but instead localize relative to the robot platform or where the sensors are connected to the soft robot. This may be acceptable if the robot is isolated from interactions with the environment. However, if the robot is meant to interact with the environment, or if it experiences any external forces (e.g. gravity loads which may cause a soft robot to deform or buckle in unexpected ways), the kinematic assumptions used with current soft robot sensors may fail.

The two strengths of outward facing cameras are first, if the cameras are able to localize accurately we will have accurate estimation that describes the configuration of our soft robot sections to which the cameras are mounted, regardless of deformations that may occur in the flexible joints. Secondly, the localization that is used to determine the configuration of the soft robot may also be used to localize relative to the environment in order to enable real-world tasks which may otherwise be difficult given uncertainty in the soft robot configuration and position relative to the environment. In this paper, we demonstrate only the first strength. However, the second strength is inherent to the localization methods that we are using for estimation.

In terms of paper organization, we first present related work in Section II. Then we present information on the robot sensing technologies that we are using in Section III. Next we present our method for soft robot configuration estimation in Section IV. Finally we present the results in Section V and conclusion in Section VI.

II. Related Work

Methods for soft robot configuration estimation range from from flex and bend sensors [2][3][4][5][6], to photodiode sensors [7], to inductance sensors [8][9], fiberoptics [10] and tactile sensors [11][12]. However, as mentioned in the introduction, all of these are sensing methods that rely on measuring the local deformation of the robot to estimate configuration (as opposed to measuring information about the environment to determine configuration and location).

Other work more closely related to that presented in this paper uses IR markers for motion capture [13][14], electromagnetic field detectors [15], or even IMUs that use the gravitational and magnetic north vector [16]. Each of these methods generally use these algorithms to estimate the orientation of some part of the soft robot. However, each of these methods (aside from the IMUs) require additional hardware that is external to the robot to estimate configuration. Unfortunately, orientation and position estimation with IMUs has the shortcoming that the estimate will drift significantly over time. To solve these problems, we turn to simultaneous localization and mapping (or SLAM) [17].

SLAM-based methods use many different sensors, but most often use cameras to develop vision-based maps and

* - equally contributing authors.

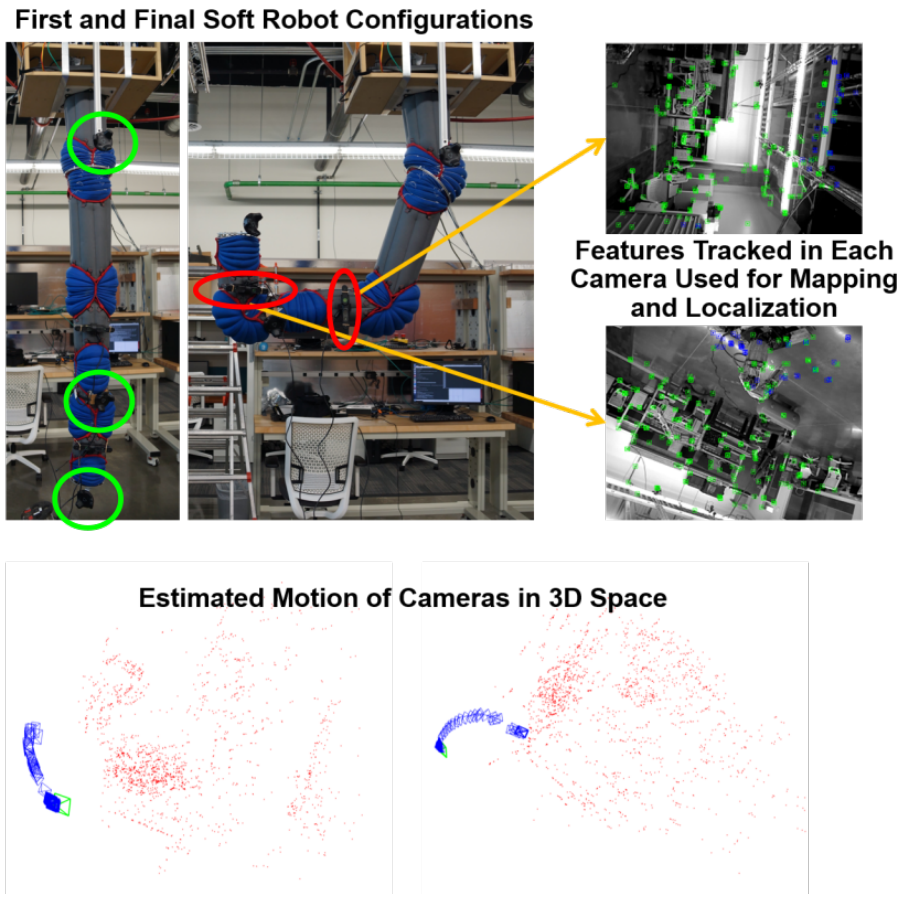

Fig. 1: Example setup of cameras running ORB-SLAM2 algorithm on the soft robot Kaa. Vive Trackers are circled in green, cameras in red. The arrows show the views corresponding to each camera. In those views, features are identified by green rectangles. The red point clouds below represent all identified features as the robot moved from the first to final configuration. These features are ORB, or Oriented FAST and Rotated BRIEF, features. More information on the SLAM algorithm and feature identification can be found in [1].

localize the camera within that map. Although not SLAM-based, similar camera-based estimation schemes have been used by looking at the inside of a soft link (see [18]) to control vibration, or by looking at the soft robot from an external perspective, requiring a camera external to the system to estimate configuration [19][20]. While in [21] they use a single camera to perform visual servoing in conjunction with a Fiber Bragg gratings (FBG) sensor rather than attempting to estimate the robot configuration from camera-based sensing.

SLAM-based methods have previously been used to estimate human motion such as in [22]. However, as far as we can tell, the efforts to estimate robot arm configuration using SLAM are limited to [23] and [24]. The first of these, [23], used a single camera mounted to a rigid manipulator to estimate joint configuration. This study is the closest to ours, as the estimation happens in real-time, however their technique was focused on estimating error between reported joint encoder values and the orientation reported by the SLAM algorithm. The second of these papers, [24], used a camera mounted at the tip of a continuum tendon-actuated robot to gather data and to predict configuration offline.

The method presented in this paper builds on these successes as we estimate the joint configuration of the robot in real time in order to enable closed-loop control of a multi-joint soft-robot, which as far as the authors know, is novel to the field.

SLAM has become a fairly mature field despite still being a fairly active area of research. Overview tutorials of the area can be found in [25][26]. Significant amounts of work are still being done on different map, landmark, and feature representations, along with the actual optimizations used to generate the maps with loop closures. Most importantly for this paper is that any progress in computational efficiency, speed, or accuracy should only improve our results.

III. Robot and Sensing Hardware

A. Robots

1) Soft Robot Kaa: Kaa is a two meter long, six DoF inflatable, pneumatically-actuated, fabric, serial link manipulator designed in collaboration with the startup company Pneubotics (see Figure 1 and [27]). The joints are composed of antagonistic, pneumatic actuation bladders which are pressurized to vary joint stiffness and apply joint torque.

The pressures are controlled using a low-level PID controller with poppet valves. The links are composed of pressurized bladders contained within ballistic-grade nylon fabric. The arm is not composed of any rigid structures.

For the first experiments performed on Kaa in this paper, we perform pressure control at each joint causing the arm to move, but not towards a goal configuration. This was done to validate our method of SLAM-based estimation independently from our attempt at closed-loop control. The next set of experiments show Kaa being controlled by a model-based predictive controller as described in [28], using our SLAM-based system for configuration estimation.

2) Baxter Robot: Baxter is a rigid robot with seven DoF for each arm (see Figure 2). We use Baxter to validate the method of SLAM-based joint estimation given the availability of ground truth with joint encoders. For the tests we perform in this paper we used the native API provided by Rethink Robotics to command different joint angles and read the corresponding joint angle measurements from encoders. For mounting the RGB-D cameras to Baxter, we 3D printed camera mounts that attached to the arm at specific locations. Despite our best efforts to design this mounting hardware, they also required adhesives to secure the cameras during operation. This is obviously one source of potential error in our joint angle estimation with Baxter.

Fig. 2: Baxter, a compliant, two-armed, robot with 7-DoF arms used for this work. The red rectangle on the arm represents where the camera was mounted. Because ground truth was available from the encoders, no Vive trackers were used to estimate pose.

B. Sensors and SLAM software

1) HTC Vive Trackers: For comparison with the SLAM-based estimation we perform on the soft robot Kaa, we used the HTC Vive virtual reality system. Configuration estimates from the HTC Vive sensors cannot be thought of as ground truth because the sensors themselves only report pose data similar to that obtained using the SLAM-based method. However, pose estimates given by the HTC Vive have been shown to be quite accurate when compared to high-fidelity motion capture systems [16] and the estimation data is available at a much higher rate than is currently possible through our SLAM-based method.

The HTC Vive system works similarly to a motion capture system, but requires only two external devices (two Lighthouse base stations, instead of a large number of cameras) along with integrated trackers attached to the robot links after every two actuated joints. These trackers are shown in Figure 1 attached to Kaa. The trackers sense infrared pulses from the Lighthouse units and combine this with data from their on-board IMUs to localize and estimate pose with approximately millimeter accuracy. We have previously used the HTC Vive Trackers to estimate orientations between soft robot links and then to estimate joint angles as discussed in [16]. The calculations needed to estimate joint angles given link orientations are sensor independent, therefore in this work the methods for joint angle estimation outlined in [16] are applied using the camera orientations returned from the SLAM algorithm.

2) RGB-D Cameras and SLAM Software: To make a map and localize the different soft robot links in their environment, we use either one or two Asus Xtion Pro-Live RGB-D cameras and the ROS package "openni_camera." The software utilized for simultaneous localization and mapping (SLAM) was an open source project called ORB-SLAM2 (see [1]).

ORB-SLAM2 tracks the location and depth of points in the environment to create a map of the surroundings. After making a map, as the cameras move through the environment subsets of those same points are found again in every frame. This makes localization of the cameras possible within the map that is created. Sample images from the SLAM software that we used can be seen in Figure 1.

ORB-SLAM2 returns visual odometry for each camera after having made an initial map, which is then converted into orientation and position data. We opted not to use the position data that ORB-SLAM2 returns as part of our SLAM-based estimation scheme. This was due to uncertainty in the camera's exact position along each robot link as a result of the inherent compliance and deformation that occur when mounting sensors to a fully compliant robot. We expect that incorporation of this data (by solving the soft bodied sensor mounting problem) would lead to improved estimates and could be done during the SLAM initialization phase. For the purposes of this paper we restricted our methods to using the orientation data reported by ORB-SLAM2.

IV. Methods

In general, our approach flips the idea of motion capture on its head: instead of cameras looking at the robot with reflective markers, we use cameras looking outward, mapping the world in order to estimate robot configuration with respect to the environment. This increases the portability of soft robotics compared to conventional motion capture. The advantage of this strategy is that despite relatively large camera localization error (at times on the order of a centimeter), each sensor is grounded in the same global frame (the map made by the data from the cameras), and therefore error does not accumulate along the robot linkage.

Our method can be used in two ways. First we can get two relative camera poses that are obtained using localization from a common map that is made by ORB-SLAM2. Or, given an initial camera pose before a soft robot segment starts moving and then a new pose measurement after the robot has moved, we can also calculate a relative camera pose for up to three representative kinematic degrees of freedom. If we expect to represent more than three degrees of freedom for a given soft robot segment (for a common parameterization that can be used for continuum joint soft robots, see [29], or lumped pin-joints, see [30]), then we would need to use multiple cameras.

Given two relative camera poses (that are obtained using localization from the map that is made by ORB-SLAM2), solving for the joint configurations is straightforward if the number of degrees of freedom in the kinematic representation between the two sensors that return pose (position and orientation) has six degrees-of-freedom or fewer. Building on our prior work where we showed that with relative poses (whether from IMUs, motion capture, or virtual reality technology) for a pair of locations on the soft robot arm, we could estimate the configuration of the soft robot between those two sensors (see [16]). In order to use the method detailed in that paper, we needed to set up a kinematic model for both Baxter and Kaa. Our algorithm uses the track-ik library [31] to perform inverse kinematics between two camera poses to estimate joint position based on the rotational output provided by ORB-SLAM2.

To validate our methodology we elected to first experiment on a rigid robot where we could compare the SLAM-based results with the ground truth angle measurements that encoders provide. For this purpose we used the Baxter robot.

A. Test 1: Single Camera SLAM Validation

To test the accuracy of our SLAM-based method when using only a single camera, we mounted an RGB-D sensor after Baxter's 3rd joint. This meant that the map was created by one camera and we estimated three degrees of freedom for Baxter based on the initial pose of the camera with the map initialized in the robots shoulder frame. The robot was commanded to 100 random configurations. Once the robot had come to rest at the given position, a set of data was recorded from Baxter's encoders and from the camera.

B. Test 2: Soft Robot Joint Estimation

Our next test included attaching an HTC Vive tracker and an RGB-D camera to the end of the large-scale, pneumatically-actuated, soft robot Kaa. Because there are not encoders or other straightforward ways to measure the actual configuration of Kaa for ground truth comparisons, we relied on previous technology that we had used for soft robot configuration estimation (see [16]). This means that all we can claim in terms of accuracy is how well the two methods agree. We do know however, that the Vive system performs comparably to high fidelity motion capture systems. So in a sense our comparison is between outward facing cameras, and traditional motion capture technology.

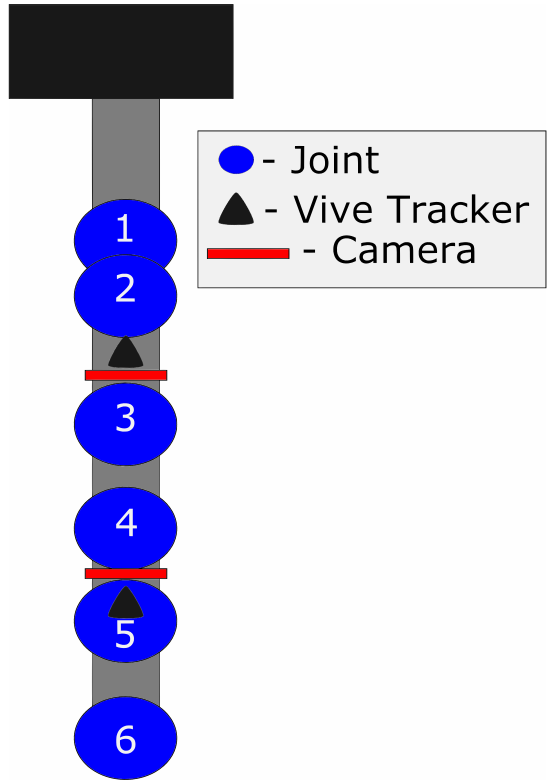

Fig. 3: Simplified diagram describing robot setup. The black box at the top of the robot is its base, and the joints are labeled 1-6 proximal to distal. Grey rectangles represent the robot's links.

We commanded the first four joints to a high pressure that would keep the first four joints pointed approximately downwards. Then, we initialized both the camera initial frame and the HTC Vive tracker frame with the last two joints facing down. Finally, we sent a randomized set of pressures to the last two joints to cause the system to move around while also recording the estimated orientations from both the Vive and the ORB-SLAM2 algorithm. A video of this trial can be seen at `https://youtu.be/vj3F4VxFhLo`. It is clear from the video that although the first few joints were commanded at constant pressures, some deformation does occur as the most distal two joints are actuated to different pressures.

C. Test 3: 4 Joint Soft Robot Control using Multi-Camera estimation

For this test we attached a camera to the first and second link of the soft robot Kaa, (see Figure 3 for a simplified description of the experimental setup). Then, after creating a map of the environment and initializing the orientation of the robot (taking great care during the map initialization that the 2 cameras aligned with the base frame of the robot and each other) we commanded the robot to move with an oscillating square wave with a period of 80s. We chose to use this length of time to show that the while the controller was not perfectly tuned, the error in the joint angles approached

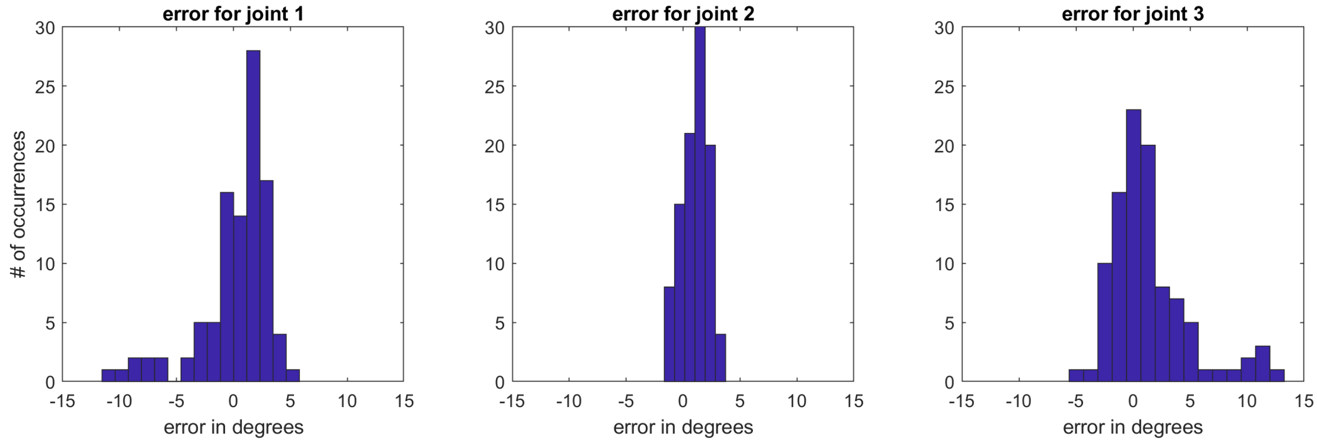

Fig. 4: Histograms showing the error between the camera estimated orientation and the ground truth orientation from the encoders on the Baxter Robot. The error shown rarely exceeds 10 degrees.

0. We also attached and recorded data from Vive trackers to compare the results between the SLAM estimation and Vive estimation (but the data from the Vive estimation was not used for control). The controller used for this test is described in [28].

V. RESULTS

A. Test 1: Single Camera SLAM Validation on Baxter Robot

For this test, the median estimation error (using the encoders as ground truth) was 1.2°, 1.3°, and 0.64° for joints 1, 2, and 3 respectively. The maximum error for joint 1, 2, and 3 was 11.5°, 3.7°, and 13.3°. Histograms of the error for the 100 different configurations are included in Figure 4. It is clear to see that the outlier errors in joint 3 are opposite the outliers for joint 1 and close to equal in magnitude. All of these errors occurred when joint 2 was close to its limit, which aligned the axis of rotation for joints 1 and 3. This may indicate that the problem is not with our sensor, but with the way we are solving the inverse kinematics problem when the axis of rotation for joint 1 and 3 are close to being parallel.

These results showed that the overall approach was viable, and could be improved by fixing some of the following sources of error. First and foremost is the initial alignment of the camera to Baxter's base frame, and to each other. Any error between the camera frame and Baxter's base frame led to poor comparison with ground truth data, and misalignment between additional cameras leads to even larger errors. This dilemma of accurately calibrating and mounting cameras is even more difficult for soft robots, and is not addressed in this paper.

Second, for this trial the camera was attached to Baxter with a prototype fixture, which made the connection between Baxter and the camera flexible. That means in some configurations the camera would be in a slightly different position or orientation on Baxter based on the direction gravity is acting on the camera in relation to Baxter or other changes in orientation due to previous movement.

Lastly, ORB-SLAM2 would lose visual features in the environment because it would get too close to the main body of the robot or otherwise have the camera view obscured, but it would recover after the obstruction passed because it could re-localize within the map it created of its environment. This last issue is one that will persist for any implementation of camera-based estimation. However, it may be mitigated with further sensor fusion such as including data from IMU sensors in an estimation scheme.

B. Test 2: Soft Robot Joint Estimation

For the second test on the soft robot Kaa, we can only report on the discrepancy between the joint angle estimation using the RGB-D camera with ORB-SLAM2 and the estimation using the Vive tracker (since there is no straightforward method to measure ground truth). A plot of the joint estimates from both sensors is shown in Figure 5.

The data for the Vive is available at 1 KHz, while the data for this specific implementation of SLAM, with this camera, and the computer we used is only updated at 10 Hz. We therefore sub-sampled the data from the Vive to line up in time with our camera-based estimates. After doing this, we could estimate the absolute value of the discrepancy between the sensors. It is important to note at this point that although we have quantified to some extent the error from the Vive estimation in the past [16], this was based on end effector position, and not on actual joint angle error. So in this case, we can not consider the Vive as ground truth, only as a benchmark with which we can compare our new estimation method.

For the second test, the median absolute discrepancy between the two methods for the first and second joints was 0.70° and 0.90°. Histograms of the difference for the 2,731 different measurements are included in Figure 6.

C. Test 3: Soft Robot Control

In this final test we again are comparing SLAM-based estimation to Vive-based estimation on Kaa, but this time instead of simply commanding pressures, we implemented a

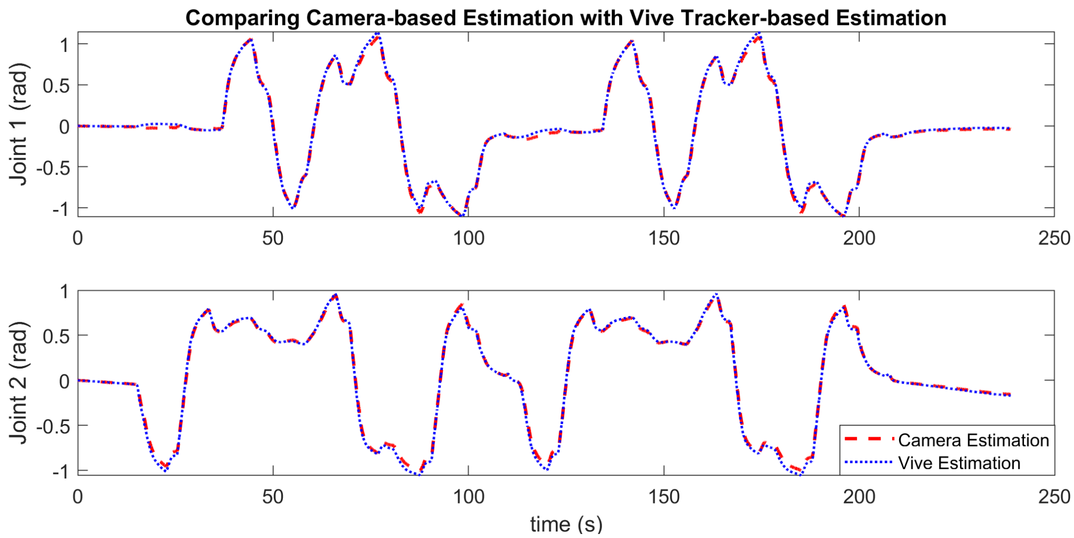

Fig. 5: Camera based versus Vive tracker estimation for the last two joints of Kaa (which we call joint 1 and joint 2 for this test). The camera based estimation follows the movement estimated using the HTC Vive trackers.

closed loop joint angle control scheme (see [28]). A plot of the control performance can be seen in Figure 7.

The data clearly show that using our SLAM based estimation method, control of a multi-degree of freedom soft robot is feasible. The SLAM estimated angle tracks almost exactly the Vive estimated angle on 3 of the 4 joints. The 3rd joint shows a discrepancy, however this was due in large part to the method that was used to fix the camera to Kaa. We used 3D printed fixtures that could be slipped over the robot when it was not inflated, and when the robot was inflated the cameras should be fixed in place because of the interference between the soft robot links and the rigid 3d print. In practice, the fixtures still allowed some motion of the cameras. In particular the fixture holding the second camera was especially susceptible in one degree of freedom because of physical constraints in mounting the camera relative to the soft robot geometry. The other issue came from our mounting point for the second Vive tracker. Because the camera was occupying the space we normally used for the Vive trackers, we connected it a little lower on the robot, which placed it on Kaa's fifth joint. This meant that due to gravity the Vive tracker would rotate at the edges of the joint trajectory, this unanticipated rotation led to the estimation scheme used for the Vive tracker to report an angle that was less then the actual angle. This manifests as the offset between the 2 state estimation techniques.

D. Discussion

The purpose of performing this research was to prove feasibility of using camera-based SLAM for configuration estimation of a soft robot. The overall results we saw were very promising. Some of the error we see is due to the frame rate of the cameras and implementation for ORB-SLAM2. Using this method with more efficient SLAM software and cameras with a faster frame rate, we believe we could reduce the error further. Additionally, implementing these measurements in conjunction with IMU measurements and a dynamic model for an estimation scheme would likely improve dynamic performance of the joint angle estimation. The most apparent weakness that we saw in our testing is the difficulty of mounting and calibrating sensors with respect to a deformable body. Because of this difficulty at the edges of our tested actuation, we saw discrepencies in our estimation. We elected to keep this data in in order to demonstrate this difficulty. Despite these weaknesses our final tests show that our SLAM-based algorithm can be used to accurately estimate joint configuration at a rate that can be used for online control of a soft robot.

One limitation of visual SLAM is that it requires sufficient visual features to generate accurate orientation measurements. This means that in environments without sufficient features to track, we would need to either track 3D features, or to augment the data from SLAM with traditional orientation sensors that can give data at a higher rate, but tend to drift or be less accurate (e.g. IMUs, see [16]). However, our results show that for a normal indoor environment (our lab) SLAM works well.

VI. Conclusion

We have demonstrated that RGB-D cameras can be used with SLAM-based localization to estimate soft robot configuration. In this paper we have raised many interesting questions that remain open (e.g. how to mount a rigid sensor to a soft robot and how to calibrate a sensor relative to the soft robot body once mounted). Without the initial evidence of feasibility in this paper, pursuing these questions and solutions further would have been unreasonable. However, given our results, we fully expect that continued research should be pursued that includes integrating cameras into soft robot fabrication and design as part of a larger effort to solve the soft robot estimation problem and enable real-world applications.

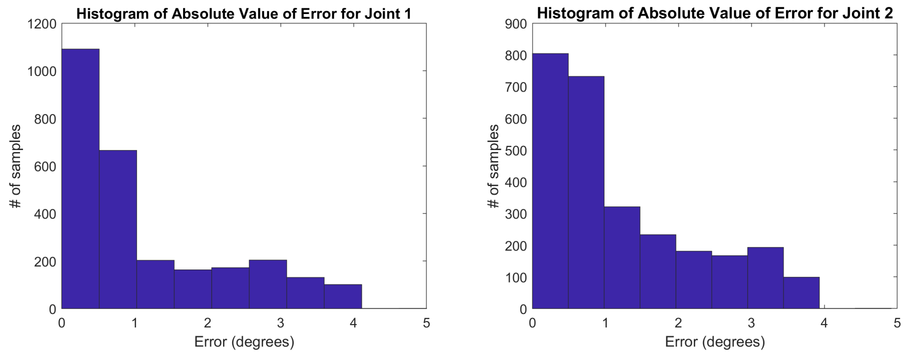

(a) (b)

Fig. 6: Histograms showing the absolute value of the discrepancy between the camera-based estimation and HTC Vive tracker-based estimation for joint 1 (a) and joint 2 (b).

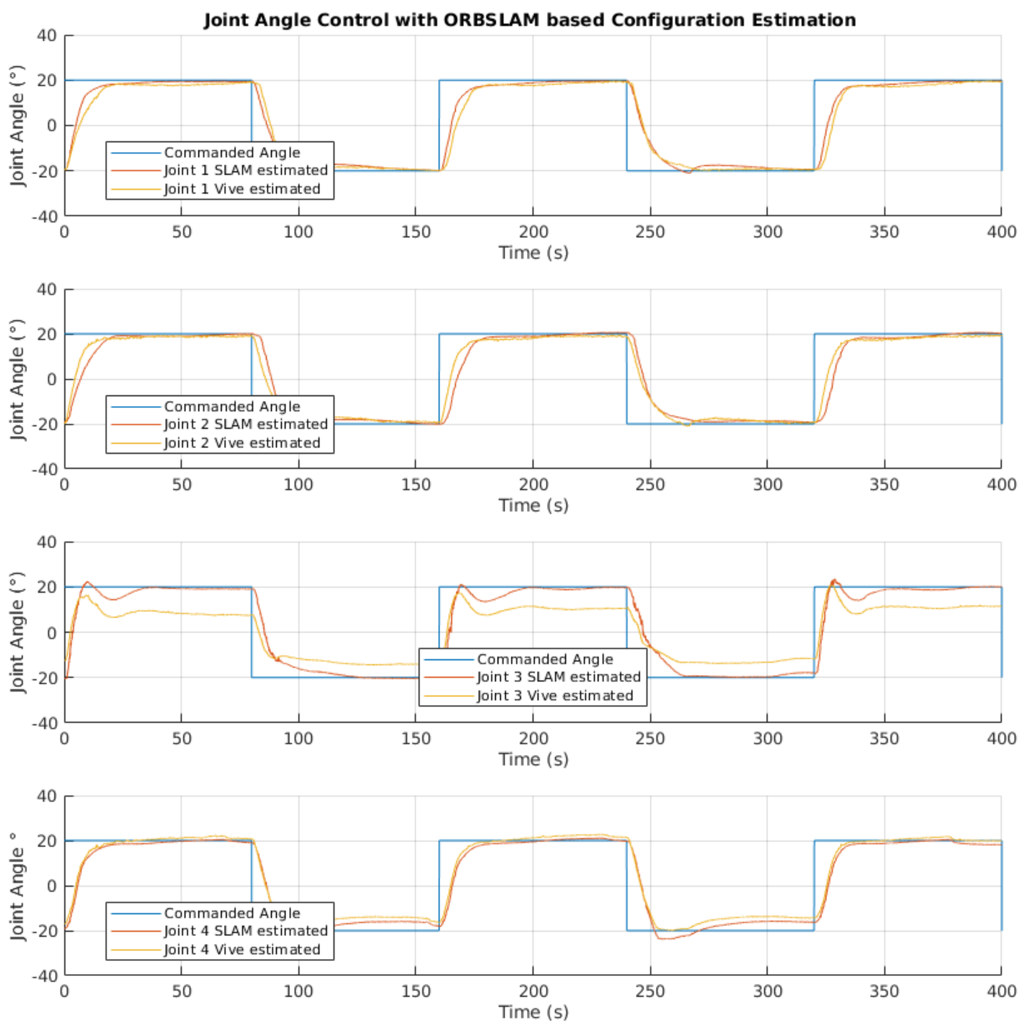

Fig. 7: Plots of each of the 4 joint angle trajectory commanded to Kaa via the MPC-based algorithm. The plots contain 2 datasets, 1 which shows the angles estimated by the SLAM algorithm, and the other showing the angles estimated by the Vive. The controller only had access to the SLAM estimated joint angles. The median tip speed during these actions was 6.2 cm/s.

Acknowledgement

We gratefully acknowledge support from NASA ECF grant NNX14A051G which has made this research possible, along with material based upon work supported by the National Science Foundation under Grant no. 1935312.

We would also like to acknowledge the efforts of Kyle Stewart in lab infrastructure development that was relevant to this paper.

References

[1] R. Mur-Artal and J. D. Tardós, "ORB-SLAM2: an open-source SLAM system for monocular, stereo and RGB-D cameras," *IEEE Transactions on Robotics*, vol. 33, no. 5, pp. 1255–1262, 2017.

[2] M. C. Yuen, H. Tonoyan, E. L. White, M. Telleria, and R. K. Kramer, "Fabric sensory sleeves for soft robot state estimation," in *2017 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, may 2017, pp. 5511–5518. [Online]. Available: http://ieeexplore.ieee.org/document/7989649/

[3] Y. She, C. Li, J. Cleary, and H.-J. Su, "Design and Fabrication of a Soft Robotic Hand With Embedded Actuators and Sensors," *Journal of Mechanisms and Robotics*, vol. 7, no. 2, p. 021007, may 2015. [Online]. Available: http://mechanismsrobotics.asmedigitalcollection.asme.org/article.aspx?doi=10.1115/1.4029497

[4] K. Elgeneidy, N. Lohse, and M. Jackson, "Data-Driven Bending Angle Prediction of Soft Pneumatic Actuators with Embedded Flex Sensors," *IFAC-PapersOnLine*, vol. 49, no. 21, pp. 513–520, 2016.

[5] P. T. Gibbs and H. H. Asada, "Wearable Conductive Fiber Sensors for Multi-Axis Human Joint Angle Measurements," *Journal of NeuroEngineering and Rehabilitation*, vol. 2, no. 1, p. 7, 2005. [Online]. Available: http://jneuroengrehab.biomedcentral.com/articles/10.1186/1743-0003-2-7

[6] T. G. Thuruthel, B. Shih, C. Laschi, and M. T. Tolley, "Soft robot perception using embedded soft sensors and recurrent neural networks," *Science Robotics*, vol. 4, no. 26, 2019. [Online]. Available: https://robotics.sciencemag.org/content/4/26/eaav1488

[7] M. K. Dobrzynski, R. Pericet-Camara, and D. Floreano, "Contactless deflection sensor for soft robots," in *IEEE International Conference on Intelligent Robots and Systems*. IEEE, sep 2011, pp. 1913–1918. [Online]. Available: http://ieeexplore.ieee.org/document/6094845/

[8] W. Felt, M. Suen, and C. D. Remy, "Sensing the motion of bellows through changes in mutual inductance," in *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, oct 2016, pp. 5252–5257. [Online]. Available: http://ieeexplore.ieee.org/document/7759772/

[9] W. Felt, M. J. Telleria, T. F. Allen, G. Hein, J. B. Pompa, K. Albert, and C. D. Remy, "An inductance-based sensing system for bellows-driven continuum joints in soft robots," *Autonomous robots*, vol. 43, no. 2, pp. 435–448, 2019.

[10] Y. Yuan, G. Wu, X. Li, Y. Fan, and X. Wu, "Effects of twisting and bending on LP_21 mode propagation in optical fiber," *Optics Letters*, vol. 36, no. 21, p. 4248, 2011. [Online]. Available: https://www.osapublishing.org/abstract.cfm?URI=ol-36-21-4248

[11] B. Shih, D. Shah, J. Li, T. G. Thuruthel, Y.-L. Park, F. Iida, Z. Bao, R. Kramer-Bottiglio, and M. T. Tolley, "Electronic skins and machine learning for intelligent soft robots," *Science Robotics*, vol. 5, no. 41, 2020. [Online]. Available: https://robotics.sciencemag.org/content/5/41/eaaz9239

[12] A. Alspach, K. Hashimoto, N. Kuppuswamy, and R. Tedrake, "Soft-bubble: A highly compliant dense geometry tactile sensor for robot manipulation," *CoRR*, vol. abs/1904.02252, 2019. [Online]. Available: http://arxiv.org/abs/1904.02252

[13] A. D. Marchese, R. Tedrake, and D. Rus, "Dynamics and trajectory optimization for a soft spatial fluidic elastomer manipulator," *The International Journal of Robotics Research*, vol. 35, no. 8, pp. 1000–1019, 2016.

[14] C. M. Best, M. T. Gillespie, P. Hyatt, L. Rupert, V. Sherrod, and M. D. Killpack, "A new soft robot control method: Using model predictive control for a pneumatically actuated humanoid," *IEEE Robotics & Automation Magazine*, vol. 23, no. 3, pp. 75–84, 2016.

[15] G. Gerboni, A. Diodato, G. Ciuti, M. Cianchetti, and A. Menciassi, "Feedback Control of Soft Robot Actuators via Commercial Flex Bend Sensors," *IEEE/ASME Transactions on Mechatronics*, vol. 22, no. 4, pp. 1881–1888, 2017.

[16] P. Hyatt, D. Kraus, V. Sherrod, L. Rupert, N. Day, and M. D. Killpack, "Configuration estimation for accurate position control of large-scale soft robots," *IEEE/ASME Transactions on Mechatronics*, vol. 24, no. 1, pp. 88–99, 2018.

[17] S. Thrun, W. Burgard, and D. Fox, *Probabilistic robotics*. MIT press, 2005.

[18] J. Oliveira, A. Ferreira, and J. C. Reis, "Design and experiments on an inflatable link robot with a built-in vision sensor," *Mechatronics*, vol. 65, p. 102305, Feb. 2020.

[19] A. Borum, D. Matthews, and T. Bretl, "State estimation and tracking of deforming planar elastic rods," in *2014 IEEE International Conference on Robotics and Automation (ICRA)*, May 2014, pp. 4127–4132.

[20] R. Wang, S. Wang, S. Du, E. Xiao, W. Yuan, and C. Feng, "Real-Time Soft Body 3D Proprioception via Deep Vision-Based Sensing," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 3382–3389, Apr. 2020.

[21] X. Wang, G. Fang, K. Wang, X. Xie, K.-H. Lee, J. D. L. Ho, W. L. Tang, J. Lam, and K.-W. Kwok, "Eye-in-Hand Visual Servoing Enhanced With Sparse Strain Measurement for Soft Continuum Robots," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 2161–2168, Apr. 2020.

[22] T. Shiratori, H. S. Park, L. Sigal, Y. Sheikh, and J. K. Hodgins, "Motion capture from body-mounted cameras," *ACM Trans. Graph.*, vol. 30, no. 4, July 2011. [Online]. Available: https://doi.org/10.1145/2010324.1964926

[23] M. Klingensmith, S. S. Sirinivasa, and M. Kaess, "Articulated robot motion for simultaneous localization and mapping (arm-slam)," *IEEE Robotics and Automation Letters*, vol. 1, no. 2, pp. 1156–1163, 2016.

[24] H. Cheng, H. Liu, X. Wang, and B. Liang, "Approximate piecewise constant curvature equivalent model and their application to continuum robot configuration estimation," in *2020 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, 2020, pp. 1929–1936.

[25] H. Durrant-Whyte and T. Bailey, "Simultaneous localization and mapping: part i," *IEEE Robotics Automation Magazine*, vol. 13, no. 2, pp. 99–110, June 2006.

[26] T. Bailey and H. Durrant-Whyte, "Simultaneous localization and mapping (slam): part ii," *IEEE Robotics Automation Magazine*, vol. 13, no. 3, pp. 108–117, Sep. 2006.

[27] D. M. Bodily, T. F. Allen, and M. D. Killpack, "Multi-objective design optimization of a soft, pneumatic robot," in *Robotics and Automation (ICRA), 2017 IEEE International Conference on*. IEEE, 2017, pp. 1864–1871.

[28] J. S. Terry, J. Whitaker, R. W. Beard, and M. D. Killpack, "Adaptive control of large-scale soft robot manipulators with unknown payloads," Oct 2019, volume 3, Rapid Fire Interactive Presentations: Advances in Control Systems; Advances in Robotics and Mechatronics; Automotive and Transportation Systems; Motion Planning and Trajectory Tracking; Soft Mechatronic Actuators and Sensors; Unmanned Ground and Aerial Vehicles. [Online]. Available: https://doi.org/10.1115/DSCC2019-9037

[29] D. M. Bodily, "Design optimization and motion planning for pneumatically-actuated manipulators," 2017.

[30] C. Della Santina, R. K. Katzschmann, A. Biechi, and D. Rus, "Dynamic control of soft robots interacting with the environment," in *2018 IEEE International Conference on Soft Robotics (RoboSoft)*. IEEE, 2018, pp. 46–53.

[31] P. Beeson and B. Ames, "Trac-ik: An open-source library for improved solving of generic inverse kinematics," in *2015 IEEE-RAS 15th International Conference on Humanoid Robots (Humanoids)*. IEEE, 2015, pp. 928–935.