

Towards Autonomous Inspection of Concrete Deterioration in Sewers with Legged Robots

Hendrik Kolvenbach¹, David Wisth², Russell Buchanan², Giorgio Valsecchi¹,
Ruben Grandia¹, Maurice Fallon², Marco Hutter¹

¹ Robotic Systems Lab
ETH Zürich, Switzerland
(hendrikk, vgiorgio, rgrandia, mahutter)@ethz.ch

² Oxford Robotics Institute
University of Oxford, United Kingdom
(davidw, russell, mfallon)@robots.ox.ac.uk

Abstract

The regular inspection of sewer systems is essential to assess the level of degradation and to plan maintenance work. Currently, human inspectors must walk through sewers and use their sense of touch to inspect the roughness of the floor and check for cracks. The sense of touch is used since the floor is often covered by (waste) water and biofilm, which renders visual inspection very challenging. In this paper, we demonstrate a robotic inspection system which evaluates concrete deterioration using tactile interaction. We deployed the quadruped robot ANYmal in the sewers of Zurich and commanded it using shared autonomy for several such missions. The inspection itself is realized via a well-defined scratching motion using one of the limbs on the sewer floor. Inertial and force/torque sensors embedded within specially designed feet captured the resulting vibrations. A pre-trained support vector machine is evaluated to assess the state of the concrete. The results of the classification are then displayed in a 3D map recorded by the robot for easy visualization and assessment. To train the SVM we recorded 625 samples with ground truth labels provided by professional sewer inspectors. We make this dataset publicly available. We achieved deterioration level estimates within three classes of more than 92% accuracy. During the four deployment missions, we covered a total distance of 300 m and acquired 130 inspection samples.

1 Introduction

The Swiss sewage systems have a total length of over 130.000 km and represent a significant communal investment. Maintaining the system is crucial for public health, but also imposes a considerable cost upon the various municipalities. Reports from industrial countries indicate that most sewage systems have reached half of the average expected lifetime of 80 years [Berger and Falk, 2009]. Thus, the focus is shifting towards maintenance and renovation. As a first step towards maintenance, cost-effective and accurate inspection is essential to assess the state of the sewer system. Sewers have been built successively, with parts of the system being over 100 years old. A large variety materials, components, and standards are used across this time. Categorization can be made between medium and large sewers, with an inner diameter of more than 800 mm and small sewers, with a diameter down to 100-150 mm. Large sewers are typically made of concrete or masonry and account for around 10% of the network [Dyk and Lohaus, 1997].

Small sewers are typically made of concrete, stoneware or plastics and account for the remaining 90%. Concrete is used in almost 40% of all sewers [Berger et al., 2016]. Although the vast majority of small sewers are of a circular or oval shape, most of the larger ones are jaw-shaped, rectangular or of irregular shape.

The rate of deterioration varies and depends on the design and usage of the sewers. Relevant factors are flow rate, slope, wastewater composition, cleaning intervals, and more [Parande et al., 2006]. An omnipresent disintegration mechanism acting on concrete in sewers is microbial induced corrosion (MIC) [Hudon et al., 2011]. This type of corrosion occurs when sulfate-reducing bacteria found in the biofilm produce hydrogen sulfide from wastewater, which is absorbed on the moist surfaces of the sewers and creates sulfuric acids. These acids react with the alkaline minerals in the concrete, which leads to the creation of large, expansive minerals and ultimately the loss of structural integrity [Wells and Melchers, 2015].

Current inspection approaches vary depending on the diameter, material, shape and expected damage to the sewer. For example, small pipes are more prone to experience clogging or leaking than large sewers. Thus, visual inspection is performed by tethered pipe-inspection robots, which crawl through the sewer and sometimes carry tools for removing blockages. Many types of small-scale robots are now commercially available [Mirats Tur and Garthwaite, 2010].

Up to now, medium to large sewers are inspected by humans with the goal of manually assessing the deterioration level. Inspectors check the roughness of the concrete visually and tactilely with their hands and feet. As the highest deterioration occurs in the center of the sewer, which is often covered by a biofilm and wastewater, purely visual assessment fails to predict the deterioration reliably. Humans are also able to adapt to irregular sewer shapes and can move through pipe diameters down to 800 mm. However, inspection tasks in typical environmental conditions (slippery ground, flowing water, dirty, damp, occasionally narrow spaces) are monotonous, dangerous and carry health-risks [Berger et al., 2016]. Additionally, extensive safety precautions are necessary such as gas sensors and safety harnesses. It is common to close off large areas of a sewage system before inspections, which potentially disrupts the network. Overall, sewer inspection presents an excellent opportunity for versatile service robots.

Only a few inspection robots for medium to large diameter sewers have been developed. *SVM-RS* from the Fraunhofer Institute for Factory Operation and Automation is a combined cleaning and inspection robot [Walter et al., 2012]. The robot weighs 2000 kg, has a size of 3500 x 1500 x 1500 mm (L x W x H) and a reach of 1200 m. Its payload includes cameras, ultrasonic sensors, structured light line scanners, and temperature sensors. Redzone Robotics' *Responder*¹ is a 300 kg, tracked platform which can be deployed in sewers with a diameter from 915-6000 mm and has a reach of up to 1000 m. The robot is equipped with cameras, laser sensors and ultrasonic sensors. A robot developed by Nanyang Technological University can be passed through tunnels with a minimum diameter of 3000 mm and has explored distances of up to 400 m [Seet et al., 2018]. The robot uses cameras and a laser profiler.

The typical method for detecting concrete damage with a robot is through an operator interpreting the acquired camera images and ultrasonic or laser sensor data. All robots use a power cable to increase their operational range. For some of the systems, deployment through a common utility hole would be difficult. An interesting robot in this context is Pure Technology's *SmartBall*², which is a free-swimming robot that can be deployed in a water stream and scans the sewer with ultrasonic sensors to find leaks. However, there is no means of controlling its path and depending on the sewer shape the deployment and recovery might be difficult.

We propose the usage of autonomous legged robots to inspect and map large and medium-sized sewers. Legged robots are relatively small, have high mobility in complex, human-made environments and can adapt their posture to inspect areas of interest [Hutter et al., 2017a] [Bellicoso et al., 2018]. Similar to humans, legged robots can probe their environment tactilely by using their limbs [Hoepflinger et al., 2010] [Kolvenbach et al., 2019]. We have successfully deployed *ANYmal* [Hutter et al., 2017b], an autonomous quadruped robot, in the sewers of Zurich as shown in previous work [Kolvenbach et al., 2019]. With the help of specially designed sensor-equipped feet, we collected a large dataset by performing an inspection motion with one limb. Later on, we trained a model to assess the level of

¹ Available at: <https://www.redzone.com/technology/responder>, Accessed: 2020-03-25

² Available at: <https://puretechltd.com/technology/smartball-leak-detection/>, Accessed: 2020-03-25

concrete deterioration using supervised machine learning techniques. In this work, we show how the trained model is used to assess the state of the concrete while performing subsequent inspection missions in the sewer³. During the missions, the robot was commanded by an operator outside the sewer using a shared autonomy framework. With this, the robot would map its environment, walk, and perform the inspections while the operator sets the next way-points and actions, respectively. A 3D map of the sewer including the inspection location and state of the concrete are the outcome of the missions. The map can be used by humans to return for repair work or monitor the deterioration rate over time.

This paper is structured as follows. First, we describe the robot with a focus on the foot design in Sec. 2. Then we present the robot deployment in the sewer (Sec. 3.1), the tactile inspection motion (Sec. 3.2), the creation of the dataset (Sec. 3.3), the classification approach (Sec. 3.4) and the mission planning system (Sec. 3.5). The outcome of the inspection missions are presented and discussed in Sec. 4. Finally, we conclude the work in Sec. 5.

2 Hardware description

2.1 ANYmal

ANYmal (Figure 1) is a 30 kg quadruped robot driven by twelve series elastic actuators mounted at the joints. The dimensions of the robot are 800 mm x 600 mm x 700 mm when standing. With its legs tucked up, it is 800 mm x 600 mm x 400 mm which allows it to be deployed through a common utility hole. The kinematic structure of the robot is designed to achieve an extensive range of motion, allowing it to overcome obstacles and manipulate the environment. The structure is made primarily of aluminum and carbon fiber to minimize weight. Shock absorbers and Kevlar plates protect exposed components from impacts. *ANYmal* is designed in a modular fashion: each actuator unit consists of motor, gear, and drive electronics and is connected over the EtherCAT bus. The connection between sub-assemblies such as the legs are easily accessible, and water-proof quick connectors are used for power and communication lines. This allows component-level ingress protection as well as fast and simple maintenance in case of hardware failure.




Figure 1: Depiction of *ANYmal* with custom feet deployed in the sewers of Zurich.

³<https://youtu.be/HjWSdmWSg>

For localization and navigation, the machine is equipped with a LIDAR (Velodyne VLP16) sensor mounted on top of the torso. A front-mounted stereo depth camera (Intel RealSense D435) provides a local map of the terrain to assess traversability and assist foothold planning. Additional sensors modules, such as RGB and thermal cameras, can be mounted for inspection tasks. Powerful LEDs provide sufficient illumination in dark environments. *ANYmal* can operate partially or fully autonomously with on-board batteries. The battery allows the system to traverse up to 3.6 km with a trotting gait at 0.5 m/s on a single charge. Optionally, the robot can be recharged without human interaction via a docking station in case long-term autonomy is required [Kolvenbach and Hutter, 2018].

2.2 Sensor-equipped, adaptive feet

We designed sensor-equipped, adaptive feet to enhance locomotion on the rough and slippery terrains encountered in sewers while measuring local ground inclination and surface properties. Similar to the rest of the robot, the feet have to be sufficiently robust to operate continuously in a challenging environment. The design is based on the adaptive foot proposed in previous work of our group [Käslin et al., 2018]. The foot consists of a large flat contact surface that can comply to the local ground inclination without interfering with the locomotion of the robot.

With a possible inclination of the terrain of up to 25° , the range of motion for ground compliance is set to 50° around the pitch axis and 30° around the roll axis. Since each leg only allows for hip abduction/adduction, hip flexion/extension and knee flexion/extension, the foot compliance around yaw prevents slipping while turning. Weighing 314 g (including cabling and connectors), it is lighter than both the original point foot and the previous adaptive foot. Figure 2a illustrates the sub-assemblies of the foot, which are described in the following.




Figure 2: Overview of the newly developed foot for sewer inspection.

1) *Foot sole*: The sole has a surface area of 60 cm^2 ($100 \text{ mm} \times 60 \text{ mm}$) and is made from an off-road rubber tire featuring 5 mm studs for increased traction (Figure 2b). The sole is connected to a metal rim by clamping, which avoids

peeling and gluing issues. Damping foam is placed between the rubber sole and the metal structure to reduce the peak loads resulting from impact forces during walking. An acetal slider avoids the foot getting stuck on overhanging edges and retains the metal rim.

2) *Pivot joint*: The pivot joint features a lightweight universal joint with integrated end stops to provide the required ground compliance. It is surrounded by an Ester Polyurethane rubber tube of Shore A70 that provides the retaining force to reset the foot to its initial position after deflection.

3) *Force sensor*: A custom, in-house developed 6-axis force/torque sensor is placed above the pivot joint to measure the forces acting on the foot. It consists of a force sensing element with strain gauges. The sensor is lightweight and robust and allows sensing up to 1000 N along the z-axis and 400 N along the x and y-axis. The maximum torque the sensor can sense is specified as 10 Nm. The accuracy lies within 1.5% of the measured value while the repeatability lies below 0.05%⁴. The measurements are filtered at 800 Hz by an integrated sinc filter with cut-off frequency of 1255 Hz. The sensor is temperature compensated to minimize drift during operation.

4) *Custom electronics*: The electronics of the foot consists of two IMUs (MPU-9250), a force sensor and a microcontroller board. One of the IMUs is located in the sole, while the other is integrated with the PCB in the shank. The IMU's linear acceleration is filtered with a cut-off frequency of 460 Hz while the angular velocity is filtered with a cut-off frequency of 184 Hz. Both the IMUs and the force sensor are connected to the microcontroller via the serial peripheral interface bus (SPI). The IMUs are read out with 1 kHz and force measurements are obtained with 400 Hz. The microcontroller board is connected to the robot via EtherCAT and powered through the auxiliary 12 V power line. The custom 6-axis force/torque features a PCB with analog-to-digital converters (ADCs) and a microcontroller that processes the analog signals of the strain gauges. Sensor data is recorded on the high-level side at 400 Hz.

5) *Shank*: The carbon fiber shank connects the foot to the knee of the robot. The shank is sealed and features a conical slider for protection of the force sensor.

6) *Sealing*: The joint is protected by thick bellows (visible in Figure 2c), mechanically clamped to the structure and sealed, which improves the ingress protection rating compared to previous work. O-rings and sealants have been used for all the matching surfaces. Water-proof cable glands and connector have been used for the cables.

3 Sewer inspection with quadruped robots

3.1 Deployment

We conducted multiple field test campaigns in the sewage system of Zurich to iterate the hardware, practice operations, collect datasets, and execute inspection missions. The first test campaign with the robot took place during two days in February 2019 [Kolvenbach et al., 2019] the second campaign took place on three days in December 2019. Generally, the robot can be easily deployed through a standard utility hole ($d = 0.8$ m). To do so, the robot's legs are unpowered, tucked up, and the system is lowered into the sewer with a tether (Figure 3). During our tests, the robot was operated from a base station which was located outside the sewer, nearby the utility hole.

The base station consists of the operator PC, an additional screen, and outlets for electricity and communication links to the robot. The base station is easy to transport and quick to set up. We installed a gazebo to shelter the robot operator and desk during rainy testing days. We deployed directional Wi-Fi antennas (Ubiquiti airMAX, HyperLink HG4958DP-19P) into the utility hole via a tripod with reversed column to enable a reliable communication link to the robot. During the test, a safety operator stayed in the sewers with a professional inspector and communicated with the robot operator outside via two-way radios (Figure 4).

The sewer inspector had to always accompany the robot for safety reasons. Potential risks include rapid changes in water levels during heavy rainfall and decreased oxygen concentration. The sewers were humid and relatively warm

⁴Available at: <https://www.botasystems.com/>, Accessed: 2020-03-25

compared to the outside temperatures during our trials in February and December 2019 because warm wastewater was flowing in the sewer. In the morning, fog was observed, which degraded visibility. Additionally, water levels would vary during the day by a few centimeters as rainfall occurred in the surrounding areas.



Figure 3: Depiction of the field test setup and the deployment of *ANYmal* into the sewer. The operator base station and warnings signs are visible in the background.

3.2 Tactile inspection motion

The scratching motion we use to inspect the floor must be repeatable and reliable across the entire range of possible surface areas. The motion, therefore, needs to be specified in such a way that it can adapt to local terrain geometry and surface roughness.

We implemented a Cartesian-space impedance controller [Khatib, 1987], which allows a motion design on both force and position level. Specifying and executing these motions was done by extending the Free Gait framework [Fankhauser et al., 2016].

The full sequence of the inspection motion, shown in Figure 5, can be split into several phases. A predefined position relative to the three stance legs is approached in (a) and contact is established in (b). In (c), a straight line trajectory is followed until a target location (d). In (e) and (f), the foot is re-positioned to return to a nominal stance on four feet. For the part (c) of the inspection motion, where data is collected, the desired end-effector force is computed as seen in the following equation.

$$\mathbf{f}_{des} = \mathbf{f}_{ref} + \mathbf{\Lambda}(\mathbf{q})\ddot{\mathbf{x}}_{ref} + \mathbf{K}_p(\mathbf{x} - \mathbf{x}_{ref}) + \mathbf{K}_d(\dot{\mathbf{x}} - \dot{\mathbf{x}}_{ref}) + c_\mu \frac{\dot{\mathbf{x}}_{ref}}{\|\dot{\mathbf{x}}_{ref}\|} |\mathbf{n}^T \mathbf{f}_{ref}|, \quad (1)$$

where \mathbf{f}_{ref} is the designed Cartesian end-effector force reference, and \mathbf{x}_{ref} , $\dot{\mathbf{x}}_{ref}$, $\ddot{\mathbf{x}}_{ref}$ are the designed position, velocity and acceleration references. $\mathbf{\Lambda}(\mathbf{q})$ is the reflected inertia matrix, which depends on the generalized coordinates \mathbf{q} . \mathbf{K}_p and \mathbf{K}_d are the position and velocity gain matrices. Finally, c_μ is a scalar value used to provide a feedforward friction compensation in the direction of the motion, scaled by the reference force along the surface normal, \mathbf{n} .

For the main part of the inspection motion, we apply a force of either 10 or 15 N on the surface. For the missions we selected 15 N. The force was experimentally tuned to keep the feet on the ground while allowing for a smooth



(a) *ANYmal* setup in the sewer



(b) Walking on biofilm



(c) Walking through wastewater

Figure 4: Depictions of *ANYmal* walking in the sewer. A professional sewer inspector accompanies the testing and assesses the state of the concrete.

scratching motion. The target location is set 100 mm forward and 50 mm sideways from the start location. Together with a total duration of 2 s and an initial and final velocity of zero, a quintic spline interpolation is defined between the start and end location to generate the motion reference. Impedance gains are set with stiffness $\mathbf{K}_p = \text{diag}(200, 200, 0)$ in N/m, and damping $\mathbf{K}_d = \text{diag}(20, 20, 20)$ in Ns/m.

Friction compensation c_μ was set to 1.0 after experimental tuning in the field. We found the high friction compensation to be important for successful motion execution on the rougher surfaces. The value of 1.0 served as a safe upper bound for the roughest terrain encountered. For more slippery surfaces the compensation is too high, but this does not pose a problem as the damping terms quickly regulate the velocity and stabilize the motion.

3.3 Dataset collection

A dataset was collected over different times in two rectangular shaped sewers to classify concrete deterioration. Both sewers were accessible through a utility hole and large enough to be traversable by humans. A slight inclination towards the center and towards the direction of flow, resulted in a higher accumulation of water in the center.

Thus, to avoid over-fitting during classification later on, we moved and reoriented the robot frequently to capture different poses of the robot with respect to the sewer floor. In total, we investigated 20 locations and areas of interest to capture a broad set of surface conditions.

Typically, the condition detection in sewers is captured by filling in evaluation sheets in which deterioration levels are grouped in classes [DIN EN 13508-2:2011, 2011]. Similar scales although more general are found in the inspection of other concrete structures, such as bridges [Everett et al., 2008].

In order to construct the dataset, we defined a scale of five condition ratings for the sewers. The scale was developed with professional sewer inspectors and is inspired by similar standards [Everett et al., 2008].

- **Good:** Smooth concrete, no problems noticeable
- **Satisfactory:** Minor signs of deterioration, increased roughness
- **Fair:** Medium signs of deterioration, increased roughness and scratches/spalling
- **Critical:** Major deterioration noticeable, large cracks, imminent failure
- **Failure:** Loss of structural integrity, leakage

The condition of the concrete we encountered in the sewers ranged from *good* to *fair*, while *critical* or extremely bad structural *failures* were not encountered. In total, we were able to collect a total of 625 samples during the two field test campaigns (*good*: 183 samples, *satisfactory*: 183 samples, *fair*: 259 samples). The samples were taken in different parts of the sewers, and were classified together with a professional sewer inspector who provided the ground truth. The dataset named *STINK* (Sewer Terrain Inspection Knowledge) is openly available⁵.

3.4 Classifying concrete deterioration

We chose a machine learning approach to capture and classify the diverse appearance of concrete deterioration together with the varying environmental conditions. As mentioned, the surface condition is not only expressed by the roughness of the concrete but also by macroscopic features such as holes, scratches, or cracks. At the same time, the surface can be dry, wet, submerged and/or covered by a biofilm (Figure 6). We extracted the raw sensor data acquired while performing the inspection motion in the sewers without further cleanup or filtering.




Figure 5: Picture sequence of the tactile inspection motion with the foot placement phase (a), inspection motion start location (b), main inspection phase (c), motion target (d), re-positioning movement (e), and final position (f).

⁵DOI: 10.3929/ethz-b-000336822 (currently being updated)

The acquired data consists of 18 signals (Force/Torque, IMU shank (linear acceleration/angular velocity), IMU sole (linear acceleration/angular velocity)). The data is cropped to 2 seconds, starting at the point of the scratching motion (Figure 5a). The raw sensor signals show a correlation between signal magnitudes and concrete deterioration level, which we exploited in the following for classification (Figure 7). The cropped sensor signals are standardized using z-scores to obtain zero mean and unit standard deviation. The dataset itself is split between randomly chosen 75% of the data for the training set (468 Samples) and 25% for the validation set (157 Samples).

Next, we decomposed the sensor signals using a fast Fourier transform (FFT) and selected the magnitude of the frequency components as features. Last, we reduced the dimensionality of the training set with a principal component analysis (PCA), which explained at least 95% of the variability. This PCA transformation of the training set was then applied to the validation set.

We utilized a support vector machine (SVM) with linear Kernel (LIBSVM for Matlab [Chang and Lin, 2011]) to train the model. The linear Kernel was chosen since it resulted in high overall accuracy. More complex, non-linear, or radial basis Kernels were omitted since they increase the chance of over-fitting while only supplying a marginal increase in performance. We performed training on the data using five-fold cross-validation combined with a grid search to find an optimal C-setting. The multi-class problem is solved by training three binary one-vs-all classifiers. The accuracy was determined by predicting the degradation level of the validation set, which was left out during training.

We were able to classify the deterioration levels with an overall classification accuracy of more than 92% and pre-




Figure 6: An exemplary set of pictures illustrating the various surface conditions encountered in the sewers.




Figure 7: Comparing the angular velocities of two samples recorded by the foot sole IMU show a correlation between signal magnitudes and concrete deterioration (Complete sequence of the inspection motion).

Table 1: Classification accuracy related to selected sensor signals (average performance over ten evaluations).

Sensor Signal	Features	Precision, Recall [%]			Accuracy
	[#]	Fair	Satisfactory	Good	[%]
Force	144	62.9, 73.3	68.2, 56.1	75.3, 69.8	67.0
Torque	171	67.5, 79.0	71.0, 64.7	77.3, 65.8	70.8
Force + Torque	214	70.9, 79.4	78.4, 69.4	79.6, 76.3	75.4
IMU Sole (linear acceleration)	297	77.1, 80.2	74.8, 70.8	86.8, 86.1	79.4
IMU Sole (angular velocity)	269	81.6, 84.0	83.1, 83.5	91.0, 85.7	84.5
IMU Sole	332	87.6, 86.5	83.4, 84.9	92.8, 92.7	87.9
IMU Shank (linear acceleration)	246	79.0, 81.7	76.1, 72.3	93.0, 92.3	82.0
IMU Shank (angular velocity)	138	72.3, 79.4	73.4, 71.0	88.2, 80.1	77.1
IMU Shank	252	85.2, 85.3	83.3, 82.1	94.4, 96.0	86.9
IMU Sole + IMU Shank	335	87.8, 90.3	88.8, 86.2	97.0, 96.1	90.7
F/T + IMU Sole + IMU Shank	341	91.6, 93.0	91.6, 89.5	95.5, 95.5	92.6

cision & recall of at least 89% on the three assigned classes. Comparing the individual sensor contributions to the classification performance shows that the combined set of sensors provides the highest accuracy (Table 1). The worst performance was found to be on the Force/Torque data, which includes the axis that is actively controlled by the robot. Nevertheless, also a smaller setup consisting of only IMU’s achieves a generally good performance. This matches previous findings on the smaller sewer inspection dataset [Kolvenbach et al., 2019] and a dataset used for planetary soil classification [Kolvenbach et al., 2019].

Performing the classification with all sensors resulted in the highest accuracy, and further investigations were performed with this selection. A confusion matrix showing the truly and wrongly classified samples was derived by training and evaluating the model for ten successive trials (Figure 8). Misclassification occurs mostly between the *fair* and *satisfactory* class, which are also fairly similar in terms of roughness compared to the smooth, *good* class. More samples were available for the *fair* class compared to the *satisfactory* and *good* class, which might explain the generally lower classification performance on the *satisfactory* class of the validation set. Investigating the sampling time shows that the signals should be captured for at least for one second for high classification accuracy. A small drop in performance is seen when taking the full sample length, which might be explained by small deviations in the duration of the scratching motion. Generally, an increase in performance might be realized by increasing the size of the dataset, further optimize the feet for vibration response, or utilizing alternative classifiers, such as neural networks [Bednarek et al., 2019].


3.5 Inspection mission

We utilized the trained SVM model to perform complete inspection missions in the sewer. The missions were planned and executed using the Director user interface (Figure 9) and system architecture, initially designed by MIT during the DARPA robotics challenge [Marion et al., 2017]. It features a shared autonomy system where the user can provide high-level commands to the robot, which are then executed autonomously. For the missions, the robot operator simply provides multiple waypoints in the online-generated map for inspection. The robot will then walk to these waypoints and perform the inspection motion autonomously.

The robot’s odometry is provided by the kinematic-inertial Two-State Implicit Filter [Bloesch et al., 2018]. This is used as a prior to a LIDAR based Iterative Closest Point (ICP) SLAM system for accurate localization and map building. During the missions, the robot also recorded a local elevation map using the forward-mounted depth sen-

		Predicted class			True class
		Fair	Satisfactory	Good	
True class	Fair	61	6.5	0.6	
	Satisfactory	6.5	39.3	1.1	
	Good	0.7	0.1	43	
		Fair	Satisfactory	Good	

(a) Confusion matrix



(b) Accuracy increase over signal duration

Figure 8: Few misclassification events occur between the *Satisfactory* and *Fair* class. Correlation between classification accuracy and signal length indicates that samples should be collected for at least a second after starting the scratching motion (Average performance over ten evaluations).

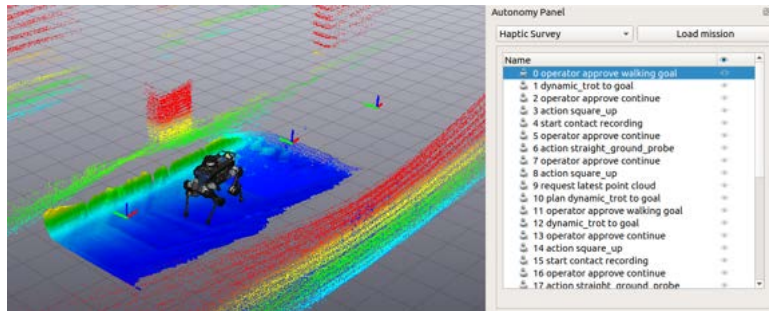


Figure 9: Screenshot of the Director mission interface [Marion et al., 2017], showing three inspection waypoints and the associated task queue.

sors [Fankhauser et al., 2018]. For each inspection motion, the position of the robot contacts in the map frame, the elevation map of the terrain around the robot, and the timestamp were recorded for post-processing.

The missions consisted of inspection walks between utility holes and tunnel crossing, which is a scenario equally performed by human inspectors. The distances between two utility holes typically range from 100-150 m, and crossing between sewers requires the robot to overcome small obstacles such as steps [Fankhauser et al., 2018]. During testing, we chose a set of different mission environments, such as a straight, featureless tunnel, a turn, and a crossing between two sewers to capture this diversity. Inspection motions were performed approximately every 3 m along the way or as commanded by the robot operator. Before each inspection motion, the robot would perform a “square up” maneuver, which would return the robot in a stable stance and consistent start position before performing the inspection.

4 Results

4.1 Mission results


Using the shared autonomy framework, we completed four missions as described in Figure 10a and Table 2. For missions 1-3, the robot was deployed through a utility hole and commanded to inspect a section of the tunnel (of length 145 m). The mission ended if either the final goal was reached or the robot slipped and could not recover. A safety operator was present at all times during the missions but did not intervene unless the robot fell. Mission 4 was designed to show the capabilities of using legged robots specifically. For this mission, the robot was placed near the window relief junction and directed by the robot operator outside the sewer to cross the window from tunnel two into tunnel one. A detailed breakdown of each mission is given below:

- **Mission 1**

The robot started in the first tunnel sewer access area and then walked down a straight section of the tunnel on the left side of the stream. This mission began in the morning when the above surface air was much cooler than the air in the tunnel. This temperature difference created fog in the sewer, as shown in Figure 10b, which was often detected as an obstacle by the LIDAR. The fog and featureless straight tunnel resulted in several artifacts in the LIDAR map (see the highlighted section of the straight tunnel in Figure 11) as well as more drift than in other missions. Since the robot walked from one utility hole to another, this mission imitates a typical inspection excursion for a human. Once the mission was complete, the robot operator directed the robot to cross the stream and walk back to the start without any interaction from the safety operator.

- **Mission 2**

The robot started in the second tunnel sewer access area and then walked towards the curved section of the



(a) Top-down diagram of the missions completed in the sewer.



(b) The robot during mission 1 as it walks down the straight featureless tunnel. Fog can be seen in front of the robot.



(c) Mission 2 in the second tunnel. Here the sewer is more curved and there is a large pipe suspended from the ceiling.



(d) View of second tunnel in direction of the window junction during mission 3. Two large windows are visible to the robot's left.



(e) The robot performing the crossing maneuver during mission 4. The robot has climbed onto the threshold and is transiting from tunnel two to tunnel one to continue the inspection.

Figure 10: Photos taken during each robot deployment.

tunnel and the window junction. In this tunnel, a pipe was suspended from the ceiling, which, along with the curvature of the tunnel in this section, provided good features for LIDAR SLAM. After reaching the window junction, the robot operator commanded the robot to cross the stream and return to the start. After crossing the stream, the robot slipped and fell into the stream, which required operator intervention to reset the robot.

- **Mission 3**

The robot started in the second tunnel sewer access area then walked towards the curved section of the tunnel and the window junction, similar to mission 2. As the robot entered the window junction, it slipped on a smooth part of concrete. The robot’s controller responded fast enough and automatically froze the joints so that the robot remained standing in place.

- **Mission 4**

One key motivation to use legged robots for sewer inspection is their versatility. We demonstrated this by placing the robot in the second tunnel and having it walk into the window junction. The robot operator then placed waypoints such that the robot climbed over a 15 cm step from tunnel two into tunnel one and the performed eight inspections. Because a legged robot can transit from one tunnel to another, they are much more useful than a robot confined to a single tunnel per mission.

Table 2: Description of the missions performed in the sewer.

Mission	Distance [m]	No. Inspections	Duration [min]
Mission 1	145	55	31
Mission 2	74	41	25
Mission 3	43	26	13
Mission 4	32	8	10
Total	294	130	79

4.2 ICP Mapping

The 3D point cloud shown in Figure 11 combines the ICP maps generated during each of the four missions, stitched together in post-processing, and overlaid on the sewer construction drawings. It shows that the point cloud reconstruction of the sewer is both internally consistent and matches the scale and geometry of the plans.

The ICP map was most accurate around the window junction due to the strong geometric features and missions 2 and 3 generated almost identical maps in independent runs. The poorest performance was during mission 1 in the long, straight portion of the tunnel. As described in Figure 10, this portion of the tunnel is almost perfectly symmetric along the direction of travel, and there were significant artifacts in the point clouds caused by dense fog, leading to an overestimate of tunnel length by 42.8%. To analyse the quality of the ICP mapping further we compared it to a ground truth point cloud collected by tripod-based commercial 3D laser scanner (Leica BLK-360). The majority of the points in the ICP map lie within 20 cm of the ground truth as shown in Fig. 12, showing the validity of this approach.

Overall, apart from scale drift in one of the missions, the point cloud is both metrically accurate and consistent with the building plans, allowing a sewer professional to easily diagnose and localize areas of concern within the sewer system.




Figure 11: Overview of whole pointcloud map overlaid on sewer plans and satellite image (1:500 scale). Blue points represent the map generated by the robot's LIDAR and the color intensity indicates height with dark blue being the sewer floor. Concrete classifications are colorized as: Red (*Fair*), Yellow (*Satisfactory*) and Green (*Good*). In green boxes we highlight two areas of interest: the curved section where missions 2 & 3 occurred and the straight section of mission 1. In the curved section we have raised the mission 3 contact maps above the ground to improve visibility.




Figure 12: Comparison of the ground truth point cloud and the ICP map generated by the robot. The ICP map shown contains only the points that lie within 20 cm of the ground truth, showing the accuracy of the ICP map. The point clouds are colorized by height for clarity.

4.3 Degradation estimates

The outcome of the inspection mission results are projected in Figure 11. In total, we were able to take 130 samples over 294 m to which we assigned a deterioration class in post-processing. There are several sections of the sewer with different levels of concrete degradation, which matched our qualitative impression. For example, the worst conditions were found between the sewer access area and window junction, while the straight portion of the tunnel and the tunnel window junction (where the robot crossed between tunnels) are less degraded. The robot generally inspected areas close to the center of the sewer where smooth terrain is not likely to be found. Mission 2 and mission 3 covered the same area in the sewer, and the classification results reflect this. Areas that have been investigated twice show a high level of consistency in the degradation estimate. Out of 45 samples that have been collected at the same patch of 43 m, just five samples show a difference in their classification result. Of the five outliers, four showed a deviation by one class and one sample by two classes. Due to the small investigated footprint of 0.005 m^2 during the inspection, it might be that the areas are not necessarily misclassified but are rather capturing a general change in deterioration of the concrete. The high spatial resolution of the measurements combined with the location in the map make it easy to visualise clusters of good or worse concrete. Generally, several inspection walks over the same areas have increased the statistical power of this approach. A possible extension of this work would be to perform the evaluation online and increase the amount of samples taken in an area with noticeable change in degradation estimates.

An example of the raw data for contact classifications is provided in Figure 13, showing the vertical ground reaction force over three inspection motions. There is a clear pattern to the data, making it easy to distinguish between the three phases of the inspection task – walking, “square up”, and inspection motion. This shows the repeatability of the robot inspection, even during complex real-world missions.

Overall, the total mission consisted of 130 inspections of which there were a total of 9 *Good* classifications, 59 *Satisfactory*, and 62 *Fair* which were colorized green, yellow and red respectively in the contact maps.




Figure 13: The vertical component of the ground reaction force of the front right foot as measured during mission 2. Walking episodes (yellow) can be easily distinguished from the “square up” reorientation maneuver (blue) and the tactile inspection motion (red). The inspection data is extracted from the sensor stream for classification.

5 Conclusion

We have shown how a legged robot can inspect concrete deterioration in medium and large sewers. Through multiple field trials, we developed a method that allows the quadruped robot *ANYmal* to walk through the sewer system of Zurich, perform tactile inspection of the concrete, and classify the deterioration level.

To cope with the wet and slippery environment in the sewer environment, we have developed adaptive planar feet

with integrated inertial and force/torque sensors. Using an impedance-controlled scratching motion, the robot can probe the terrain with one of its limbs while maintaining balance with the others. We acquired and open-sourced a dataset with 625 samples consisting of the sensor data and ground truth labels supplied by professional sewer inspectors. Training a support vector machine on the dataset allowed us to predict the current state of the concrete deterioration within three classes with overall more than 92%, as well as more than 89% precision and recall. Analysis of the data shows that even a small set of sensors is sufficient for high classification accuracy.

We planned and executed several inspection missions with the robot, which mimic the approach of human inspectors. Using a shared autonomy framework, we were able to command the robot from outside the sewer in individual missions of up to 145 m while performing the concrete inspection action at regular intervals. While walking, the robot would map the sewer using on-board sensors and record the points of inspection. In post-processing, the state of the concrete was inspected using the pre-trained SVM model. We executed four missions in different areas of the sewer, including climbing between sewers, which demonstrates the advantage of using a legged system in this environment. Within the four missions, we covered a total distance of 294 m in 80 min and performed 130 inspection motions.

While the inspection approach worked well, further improvements are required to increase robustness and reliability. First, the system needs to be tested in narrow, partially flooded, and highly slippery sewers to demonstrate a similar capability to a human inspector. This also includes further robustifying the locomotion system against slip-induced falling. Secondly, more data needs to be collected to confirm the robustness on a more diverse set of sewers, to improve classification accuracy, and eventually, the precision of the prediction. Lastly, the mapping could be improved to cope with straight, featureless tunnels and to increase the accuracy of the inspection locations. Overall, we believe that these shortcomings will be overcome and that legged robots will become a valuable partner in the inspection of sewers.

Acknowledgments

The authors want to thank Roman Weiss and his colleagues from *ERZ - Entsorgung + Recycling Zurich* for their support during the field tests. Klajd Lika and Ilias Patsiaouras from *Bota Systems* helped to integrate the force/torque sensors into the feet and Marco Camurri built the ground truth map. This work has been supported by the European Union's Horizon 2020 research and innovation programme under grant agreement No 780883 and by the European Space Agency (ESA) and Airbus DS in the framework of the Network Partnering Initiative 481-2016. This work has been conducted as part of ANYmal Research, a community to advance legged robotics.

References

- [Bednarek et al., 2019] Bednarek, J., Bednarek, M., Wellhausen, L., Hutter, M., and Walas, K. (2019). What am i touching? learning to classify terrain via haptic sensing. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 7187–7193.
- [Bellicoso et al., 2018] Bellicoso, C. D., Bjelonic, M., Wellhausen, L., Holtmann, K., Günther, F., Tranzatto, M., Fankhauser, P., and Hutter, M. (2018). Advances in real-world applications for legged robots. *Journal of Field Robotics*, 35(8):1311–1326.
- [Berger and Falk, 2009] Berger, C. and Falk, C. (2009). *Zustand der Kanalisation in Deutschland - Ergebnisse der DWA-Umfrage 2009*. Deutsche Vereinigung für Wasserwirtschaft, Abwasser und Abfall e. V. (DWA), Hennef.
- [Berger et al., 2016] Berger, C., Falk, C., Hetzel, F., Pinnekamp, J., Roder, S., and Ruppelt, J. P. (2016). *Zustand der Kanalisation in Deutschland - Ergebnisse der DWA-Umfrage 2015*. Deutsche Vereinigung für Wasserwirtschaft, Abwasser und Abfall e. V. (DWA), Hennef.
- [Bloesch et al., 2018] Bloesch, M., Burri, M., Sommer, H., Siegwart, R., and Hutter, M. (2018). The two-state implicit filter recursive estimation for mobile robots. *IEEE Robotics and Automation Letters*, 3(1):573–580.

- [Chang and Lin, 2011] Chang, C.-C. and Lin, C.-J. (2011). LIBSVM: A library for support vector machines. *ACM Transactions on Intelligent Systems and Technology*, 2:27:1–27:27. Software available at <http://www.csie.ntu.edu.tw/~cjlin/libsvm>.
- [DIN EN 13508-2:2011, 2011] DIN EN 13508-2:2011 (2011). Investigation and assessment of drain and sewer systems outside buildings - part 2: visual inspection coding system.
- [Dyk and Lohaus, 1997] Dyk, C. and Lohaus, J. (1997). *Der Zustand der Kanalisation in Deutschland - Ergebnisse der ATV-Umfrage 1997*. ATV (Abwassertechnische Vereinigung e.V.), Hennef.
- [Everett et al., 2008] Everett, T. D., Weykamp, P., Capers, H. A., Cox, W. R., Drda, T. S., Hummel, L., Jensén, P., Juntunen, D. A., Kimball, T., and Washer, G. (2008). Bridge evaluation quality assurance in europe.
- [Fankhauser et al., 2016] Fankhauser, P., Bellicoso, C. D., Gehring, C., Dubé, R., Gawel, A., and Hutter, M. (2016). Free gait: An architecture for the versatile control of legged robots. In *2016 IEEE-RAS 16th International Conference on Humanoid Robots (Humanoids)*, pages 1052–1058.
- [Fankhauser et al., 2018] Fankhauser, P., Bjelonic, M., Dario Bellicoso, C., Miki, T., and Hutter, M. (2018). Robust rough-terrain locomotion with a quadrupedal robot. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 5761–5768.
- [Fankhauser et al., 2018] Fankhauser, P., Bloesch, M., and Hutter, M. (2018). Probabilistic terrain mapping for mobile robots with uncertain localization. *IEEE Robotics and Automation Letters (RA-L)*, 3(4):3019–3026.
- [Hoepflinger et al., 2010] Hoepflinger, M. A., Remy, C. D., Hutter, M., Spinello, L., and Siegwart, R. (2010). Haptic terrain classification for legged robots. In *2010 IEEE International Conference on Robotics and Automation*, pages 2828–2833.
- [Hudon et al., 2011] Hudon, E., Mirza, S., and Frigon, D. (2011). Biodeterioration of concrete sewer pipes: State of the art and research needs. *Journal of Pipeline Systems Engineering and Practice*, 2(2):42–52.
- [Hutter et al., 2017a] Hutter, M., Diethelm, R., Bachmann, S., Fankhauser, P., Gehring, C., Tsounis, V., Lauber, A., Guenther, F., Bjelonic, M., Isler, L., Kolvenbach, H., Meyer, K., and Hoepflinger, M. (2017a). Towards a Generic Solution for Inspection of Industrial Sites. In *Field and Service Robots (FSR)*.
- [Hutter et al., 2017b] Hutter, M., Gehring, C., Lauber, A., Gunther, F., Bellicoso, C. D., Tsounis, V., Fankhauser, P., Diethelm, R., Bachmann, S., Bloesch, M., Kolvenbach, H., Bjelonic, M., Isler, L., and Meyer, K. (2017b). Any-mal - toward legged robots for harsh environments. *Advanced Robotics*, 31(17):918–931.
- [Käslin et al., 2018] Käslin, R., Kolvenbach, H., Paez, L., Lika, K., and Hutter, M. (2018). Towards a passive adaptive planar foot with ground orientation and contact force sensing for legged robots. *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2018)*.
- [Khatib, 1987] Khatib, O. (1987). A unified approach for motion and force control of robot manipulators: The operational space formulation. *IEEE Journal on Robotics and Automation*, 3(1):43–53.
- [Kolvenbach et al., 2019] Kolvenbach, H., Bärtschi, C., Wellhausen, L., Grandia, R., and Hutter, M. (2019). Haptic inspection of planetary soils with legged robots. *IEEE Robotics and Automation Letters*, 4(2):1626–1632.
- [Kolvenbach and Hutter, 2018] Kolvenbach, H. and Hutter, M. (2018). Life extension: An autonomous docking station for recharging quadrupedal robots. In *Field and Service Robotics*, pages 545–557, Cham. Springer International Publishing.
- [Kolvenbach et al., 2019] Kolvenbach, H., Valsecchi, G., Grandia, R., Ruiz, A., Jenelten, F., and Hutter, M. (2019). Tactile inspection of concrete deterioration in sewers with legged robots. In *Field and Service Robotics (FSR)*.
- [Marion et al., 2017] Marion, P., Fallon, M., Deits, R., Valenzuela, A., D’Arpino, C. P., Izatt, G., Manuelli, L., Antone, M., Dai, H., Koolen, T., Carter, J., Kuindersma, S., and Tedrake, R. (2017). Director: A user interface designed for robot operation with shared autonomy. *Journal of Field Robotics*, 34:225–426.

- [Mirats Tur and Garthwaite, 2010] Mirats Tur, J. M. and Garthwaite, W. (2010). Robotic devices for water main in-pipe inspection: A survey. *Journal of Field Robotics*, 27(4):491–508.
- [Parande et al., 2006] Parande, A. K., Ramsamy, P. L., Ethirajan, S., Rao, C. R. K., and Palanisamy, N. (2006). Deterioration of reinforced concrete in sewer environments. *Proceedings of the Institution of Civil Engineers - Municipal Engineer*, 159(1):11–20.
- [Seet et al., 2018] Seet, G., Yeo, S., Law, W., Burhan, Wong, C., Sapari, S., and Liao, K. (2018). Design of tunnel inspection robot for large diameter sewers. *Procedia Computer Science*, 133:984 – 990. International Conference on Robotics and Smart Manufacturing (RoSMa2018).
- [Walter et al., 2012] Walter, C., Saenz, J., Elkmann, N., Althoff, H., Kutzner, S., and Stuerze, T. (2012). Design considerations of robotic system for cleaning and inspection of large-diameter sewers. *Journal of Field Robotics*, 29(1):186–214.
- [Wells and Melchers, 2015] Wells, T. and Melchers, R. (2015). Modelling concrete deterioration in sewers using theory and field observations. *Cement and Concrete Research*, 77:82 – 96.