
White Rose Research Online

eprints@whiterose.ac.uk
https://eprints.whiterose.ac.uk
Universities of Leeds, Sheffield and York

Deposited via The University of Sheffield.

White Rose Research Online URL for this paper:
https://eprints.whiterose.ac.uk/id/eprint/241148/

Version: Accepted Version

Proceedings Paper:
Karnezis, A., Worley, R., Anderson, S.R. et al. (2026) Probabilistic pipe material classification via LiDAR-IMU data fusion. In: 2026 29th International Conference on Information Fusion (FUSION). 29th International Conference on Information Fusion (FUSION 2026), 23-26 Jun 2026, Trondheim, Norway. Institute of Electrical and Electronics Engineers (IEEE). ISBN: 9798331552350. ISSN: 2707-8787. EISSN: 2707-8787.

https://doi.org/10.23919/FUSION67272.2026.11726554

© 2026 The Author(s). Except as otherwise noted, this author-accepted version of a paper published in 29th International Conference on Information Fusion (FUSION) is made available via the University of Sheffield Research Publications and Copyright Policy under the terms of the Creative Commons Attribution 4.0 International License (CC-BY 4.0), which permits unrestricted use, distribution and reproduction in any medium, provided the original work is properly cited. To view a copy of this licence, visit http://creativecommons.org/licenses/by/4.0/

Reuse
This article is distributed under the terms of the Creative Commons Attribution (CC BY) licence. This licence allows you to distribute, remix, tweak, and build upon the work, even commercially, as long as you credit the authors for the original work. More information and the full terms of the licence here: https://creativecommons.org/licenses/

Takedown
If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.





Probabilistic Pipe Material Classification via LiDAR-IMU Data Fusion

Aristeidis Karnezis
School of Electrical & Electronic Engineering
University of Sheffield
Sheffield, United Kingdom
a.karnezis@sheffield.ac.uk

Rob Worley
School of Electrical & Electronic Engineering
University of Sheffield
Sheffield, United Kingdom
r.worley@sheffield.ac.uk

Sean R. Anderson
School of Electrical & Electronic Engineering
University of Sheffield
Sheffield, United Kingdom
s.anderson@sheffield.ac.uk

Simon J. Tait
School of Mechanical, Aerospace & Civil Engineering
University of Sheffield
Sheffield, United Kingdom
s.tait@sheffield.ac.uk

Kirill V. Horoshenkov
School of Mechanical, Aerospace & Civil Engineering
University of Sheffield
Sheffield, United Kingdom
k.horoshenkov@sheffield.ac.uk

Lyudmila Mihaylova
School of Electrical & Electronic Engineering
University of Sheffield
Sheffield, United Kingdom
l.s.mihaylova@sheffield.ac.uk

***Abstract*—Reliable identification of pipe material in underground sewer networks is essential for targeted maintenance and asset management. This paper presents a probabilistic framework for classifying clay, concrete, and plastic pipes by fusing LiDAR intensity with IMU-based joint detection. LiDAR intensity is normalised for range and scan angle effects using a shared bivariate polynomial model fitted across all materials, producing residuals that are sensitive to material type and whose distributions are captured by kernel density estimation (KDE). Independently, the Inertial Measurement Unit (IMU) detects mechanical perturbations at pipe joints. The gaps between detected joints are compared with the expected segment length for each material. Long gaps can be interpreted as missed-joint cases when they are close to integer multiples of the expected segment length, while unreliable gaps are down-weighted by an outlier component. Both modalities share a unified Bayesian framework with uniform priors, and are combined in the log domain using mean log-likelihoods with an adjustable modality weight. Experiments use six controlled runs across three single-material pipe networks with distinct segment lengths. In this setting, intensity alone achieves posterior confidence of approximately 49-60%, while distance alone reaches 69-86.5%. Fusion increases posterior confidence relative to either single modality in all cases, assigning 93.5%, 93.5%, and 86.2% posterior probability to the true concrete, clay, and plastic classes, respectively.**

***Index Terms*—Pipe inspection, material classification, LiDAR intensity, IMU, Bayesian inference, data fusion**

I. Introduction

Underground sewer networks are critical components of civil infrastructure, yet their effective management is often constrained by incomplete knowledge of the assets themselves. In many systems, pipes have been installed, repaired, and extended over several decades, resulting in networks that contain a heterogeneous mixture of materials, including vitrified clay, concrete, and plastic. Identifying the material type of each pipe segment becomes a fundamental requirement for asset management, since it directly informs rehabilitation strategy selection, estimation of remaining service life, and the prioritisation of maintenance and renewal budgets [1]. In practice, however, asset and inspection records are frequently incomplete, inconsistent, or otherwise unreliable [2], a challenge that is widely acknowledged in the sewer asset management literature [3], [4].

To address these information gaps, robotic inspection has become the primary tool for sewer asset assessment [5], [6]. Inspection robots can traverse pipes that are inaccessible to human operators and collect rich streams of sensor data in GPS-denied environments. As a result, much of the robotics research in this domain has focused on localisation and mapping, typically using combinations of wheel odometry, inertial sensing, acoustics, and laser profiling to support operation in long, featureless pipe networks [7], [8]. At the same time, these robots carry sensing modalities that capture information beyond geometry and motion, creating an opportunity to estimate pipe properties such as material type directly from existing inspection data streams, without requiring additional hardware or dedicated inspection passes.

This opportunity is attractive, but material identification from inspection data remains challenging because no single sensing modality is sufficiently reliable under all operating conditions. LiDAR-equipped inspection robots, for instance, provide both range and return intensity measurements [9]. Because different pipe materials exhibit different optical reflectance characteristics at infrared laser wavelengths, raw intensity values carry useful information about material type

We are grateful for the support via the EU project: "PIPEON: Robotics and AI for Sewer Pipe Inspection and Maintenance" under grant number 101189847, the UK EPSRC Pipebots: Pervasive Sensing for Buried Pipes Programme Grant (EP/S016813/1) and Ofwat funded Pipebots Patrol project.

[10], [11]. However, the measured intensity is also strongly affected by geometric factors - specifically range and angle of incidence - as well as surface conditions including moisture, dirt, and biofilm [12]. These confounds must be addressed before intensity can reliably classify pipe materials.

A complementary source of material information arises from the physical structure of the network. When an inspection robot crosses a joint between pipe segments, the discontinuity produces a transient perturbation in the Inertial Measurement Unit (IMU) signal. By detecting these perturbations and measuring the distance between detected joints, it is possible to estimate segment length and use it as an indirect indicator of material type. Unlike LiDAR intensity, this signal is insensitive to optical surface contamination. Its weaknesses though are sparsity and dependence on reliable joint detection, since missed detections can produce longer observed gaps.

In this work, segment length is not treated as a universal rule for each material. It can vary with the manufacturer, pipe type, installation era, and local repairs. Instead, we use the reference runs to learn the segment lengths characteristic of the tested networks, and the model expresses uncertainty when this evidence is weak.

These complementary limitations motivate a fusion approach in which the two modalities compensate for one another. This paper presents a unified probabilistic Bayesian framework that combines geometrically normalised LiDAR intensity with IMU-derived distances between joints to classify sewer pipe material. The two modalities are modelled independently and fused in the log domain with adjustable weighting, producing calibrated posterior probabilities over material classes without requiring additional sensing hardware.

Three networks (one concrete, one clay, and one plastic) are used to build reference models for both sensing modalities, and three additional networks serve as held-out targets for classification. In these controlled experiments, fusion improves confidence and resolves ambiguities that persist when LiDAR intensity or IMU segment length cues are used in isolation.

The main contributions of this work are:

1) a LiDAR intensity normalisation approach based on bivariate polynomial modelling of range and scan angle effects, producing material-sensitive intensity residuals suited to the cylindrical pipe geometry,
2) an IMU-based method that detects pipe joints and compares the inter-joint gaps with expected segment lengths, while explicitly allowing a long gap to represent one or more missed joints rather than treating it as noise,
3) a log domain fusion strategy with adjustable modality weighting that produces calibrated posterior probabilities over material classes for each segment, and
4) an experimental validation across single material sewer networks demonstrating that fusion improves classification confidence relative to intensity alone and resolves ambiguities between modalities.

The rest of the paper is organised as follows. Section II reviews related work. Section III presents the proposed probabilistic classification framework, covering the shared Bayesian formulation, the LiDAR intensity and IMU distance modalities, and their fusion. Section IV describes the experimental setup. Section V presents the classification results, and Section VI discusses the findings and outlines future work.

II. Related Work

A. Robotic localisation and the role of joints

Localisation in GPS-denied, geometrically repetitive pipe networks typically relies on dead-reckoning from wheel odometry and inertial sensing, sometimes augmented with topological constraints.

Aitken et al. [13] surveyed localisation and mapping methods for pipe inspection robots across a range of sensor modalities and identified IMU odometry fusion as a robust and widely used baseline in these environments. Corominas Murtra and Mirats Tur [14] proposed a graph optimisation approach that fused IMU measurements with cable encoder length data, demonstrating robust pose estimation in a tethered sewer inspection scenario. Edwards et al. [15] exploited pipe joint spacing as a one degree of freedom localisation signal, using visual joint detection and demonstrating robustness advantages over ORB-SLAM3 in feature-sparse pipes.

Across this literature, pipe joints appear as some of the most reliable discrete features available to the robot. Because joints occur at approximately regular intervals and produce observable signatures in visual or inertial data, they have been used as landmarks for drift correction and localisation support [16], [17], but not as indicators of pipe material. The present work exploits the distance between detected joints as a material-dependent signal. To the authors' knowledge, inter-joint distance has not previously been used as a probabilistic cue for pipe material classification in the sewer robotics literature.

B. LiDAR intensity for material classification

A central finding of the broader LiDAR literature is that raw intensity cannot be interpreted directly as a material signature because it is confounded by data acquisition geometry, most importantly target range and beam incidence angle.

Höfle and Pfeifer [11] provided an early systematic treatment of these confounds and proposed correction strategies based on the radar range equation. Jutzi and Gross [10] proposed a data-driven correction model that estimates the combined effects of range and incidence angle simultaneously via log-domain linear regression, without requiring external calibration targets. Kaasalainen et al. [18] showed that normalisation is necessary before intensity values can be meaningfully compared across measurements. Kashani et al. [19] provide a comprehensive review of the field, distinguishing intensity correction, normalisation, and rigorous radiometric calibration as distinct processing levels. A relevant finding for sewer environments specifically is that moisture can substantially reduce near infrared reflectance [20], highlighting the importance of empirical correction in wet and fouled pipes.

These issues are particularly pronounced in the cylindrical sewer pipe geometry, where the sensor to wall distance and the

incidence angle vary systematically within every scan profile. The present work follows Jutzi and Gross [10] in adopting an empirically fitted joint model of range and angle, but replaces their log-linear parametric form with a bivariate polynomial that makes no assumption about the functional form of the combined dependence.

Although LiDAR intensity has been exploited for surface material classification in other domains [12], [21], its application to *pipe material classification in sewer inspection* remains limited.

C. Probabilistic sensor fusion

Combining complementary but imperfect sensor cues within a probabilistic framework is a longstanding theme in robotics and pattern recognition [22], [23]. Duran et al. [24] surveyed sensor technologies for sewer inspection and highlighted the limitations of any single modality. However, such work focuses on detecting *damage* or quantifying *condition*, not on identifying pipe material.

The fusion of LiDAR intensity residuals with IMU-derived inter-joint distances for pipe material classification remains unexplored. The framework proposed in this paper fills this gap, leveraging signals already present in standard inspection payloads without requiring new hardware.

III. Probabilistic Classification Framework

This section presents the shared Bayesian formulation that applies to any sensing modality (Section III-A), then develops the LiDAR intensity model (Section III-B), the IMU-based distance model (Section III-C), and finally the fusion strategy that combines both (Section III-D).

A. Bayesian Formulation

Both LiDAR intensity and IMU derived distances between joints are treated as independent sources of information about the same unknown quantity: the material m of the pipe segment under inspection. The framework developed here applies to either modality individually and, in Section III-D, to their combination. Let $\mathbf{x}$ denote the observations from a single source, and let $\mathcal{M} = \{\text{Concrete}, \text{Clay}, \text{Plastic}\}$ be the set of candidate materials.

1) Posterior Probability via Bayes' Rule: The posterior probability of material m given observations $\mathbf{x}$ is:

$$p(m \mid \mathbf{x}) = \frac{p(\mathbf{x} \mid m) \cdot p(m)}{\sum_{m' \in \mathcal{M}} p(\mathbf{x} \mid m') \cdot p(m')}, \tag{1}$$

where $p(\mathbf{x} \mid m)$ is the likelihood of the observations under material model m, and $p(m)$ is the prior probability of that material. In the absence of prior information about which material is present at a given location, a uniform prior is adopted. So the posterior simplifies to:

$$p(m \mid \mathbf{x}) = \frac{p(\mathbf{x} \mid m)}{\sum_{m' \in \mathcal{M}} p(\mathbf{x} \mid m')}. \tag{2}$$

The posterior is therefore driven entirely by the likelihood under each material model. If two likelihoods are similar, the posterior reflects genuine ambiguity rather than favouring one material, motivating fusion to resolve such cases.

Working directly with likelihoods $p(\mathbf{x} \mid m)$ can be numerically unstable, because these values can become very small for large observation sets. We therefore work with the log-likelihood:

$$\ell_m = \log p(\mathbf{x} \mid m). \tag{3}$$

The posterior probabilities are then computed by normalising these log-likelihoods over the candidate materials. This is done using the log-sum-exp identity:

$$\log p(m \mid \mathbf{x}) = \ell_m - \ell^* - \log \sum_{m'} \exp(\ell_{m'} - \ell^*), \tag{4}$$

where $\ell^* = \max_{m'} \ell_{m'}$. Subtracting ℓ^* before exponentiation keeps the computation numerically stable, because the largest exponent becomes $\exp(0) = 1$. The averaging over measurements is performed inside the log-likelihood for each sensing modality, as shown for LiDAR intensity in (15) and for IMU distance in (22).

2) Classification Decision and Confidence: The predicted material is the one with the highest posterior probability:

$$\hat{m} = \arg\max_{m \in \mathcal{M}} \; p(m \mid \mathbf{x}). \tag{5}$$

The confidence of the prediction is defined as the posterior probability assigned to the winning class:

$$c = \max_{m \in \mathcal{M}} \; p(m \mid \mathbf{x}). \tag{6}$$

A value of $c \approx 1$ indicates that one material is strongly favoured while a value of $c \approx 1/3$ indicates that the observations are equally consistent with all models and the classification is unreliable.

B. LiDAR Intensity Classification

LiDAR scan intensity forms a repeatable signal along the pipe that can indicate material type. The main difficulty is that the measured intensity also changes with measurement geometry, principally the distance to the wall r and the scan angle θ. With the LiDAR mounted at a fixed tilt ($30°$), range and angle vary predictably across each scan, so two points on the same material can produce different measured intensities. The method estimates how intensity changes with distance and scan angle, subtracts this expected variation, and the remaining residuals are then compared between materials.

1) Spatial Coordinate and Intensity Normalisation: Each LiDAR measurement is timestamped, and the robot moves at a known, approximately constant speed v. Timestamps are converted to a coordinate x representing distance along the pipe. Let t denote the LiDAR timestamp and t_0 the first timestamp of a run. The LiDAR reference position along the pipe is approximated as:

$$x(t) = v \; (t - t_0) + x_0, \tag{7}$$

where x_0 is a constant offset used to align the start of the measurement with the physical pipe coordinate.

A LiDAR beam typically strikes the wall slightly ahead of the sensor reference point, depending on beam angle and measured range. Ignoring this causes measurements near a material boundary to be assigned to the wrong distance bin, blurring the transition. A forward projection corrects the assigned position of each return. Let r be the measured range and θ the scan angle in radians. The projected hit coordinate becomes:

$$x_{\text{hit}}(r, \theta, t) = x(t) + \Delta x, \tag{8}$$

with:

$$\Delta x = \begin{cases} 0, & r\cos\theta \leq 0, \\ r\cos\theta, & r\cos\theta > 0. \end{cases} \tag{9}$$

All points are assigned to fixed spatial bins along the pipe. Points are filtered by range using $r \in [r_{\min}, r_{\max}]$, with $r_{\min} = 0.05\,\text{m}$ and $r_{\max} = 0.25\,\text{m}$, to keep only stable returns. No angular sector restriction is applied as the range gate alone is sufficient to exclude invalid measurements in this geometry.

2) Accounting for Geometry in LiDAR Intensity: Let $\mathcal{I}_M$ denote the measured LiDAR intensity for a valid return. Even when the material is unchanged, $\mathcal{I}_M$ varies with r and θ because measurement geometry affects return strength. The expected intensity due to geometry is modelled by a smooth function $\mathcal{I}_E(r, \theta)$, approximated as a low-order bivariate polynomial. For numerical stability, inputs are first standardised:

$$z_r = \frac{r - \mu_r}{s_r}, \qquad z_\theta = \frac{\theta - \mu_\theta}{s_\theta}, \tag{10}$$

where μ_r, μ_θ are empirical means, and s_r, s_θ are empirical standard deviations, computed over the pooled reference data. The geometric intensity model is then:

$$\mathcal{I}_E(r, \theta) = \sum_{i=0}^{D_r} \sum_{j=0}^{D_\theta} \beta_{ij}\ {z_r}^i\, {z_\theta}^j, \tag{11}$$

with degrees $D_r = D_\theta = 2$. The polynomial is fitted by least squares to the pooled dataset across all materials, using measurements between the 5th and 95th percentile of intensity to exclude outliers. Because the same geometric baseline is subtracted from every measurement regardless of material identity, the resulting residuals can be computed for unlabelled data at classification time without requiring prior knowledge of the pipe material.

The normalised residual is then:

$$\mathcal{I}_R = \mathcal{I}_M - \mathcal{I}_E(r, \theta). \tag{12}$$

A positive $\mathcal{I}_R$ indicates a stronger return than geometry predicts while a negative value indicates a weaker one. Subtracting the geometric baseline $\mathcal{I}_E(r, \theta)$ reduces the influence of range and scan angle, making residuals more comparable across viewing angles and repeated runs. This correction is purely geometric: it removes the systematic dependence on range and scan angle, but it does not directly correct for moisture, dirt, biofilm, ageing, or surface roughness. These factors continue to affect the LiDAR intensity, which helps explain why the intensity-only result can remain uncertain even after correction. For this reason, the fusion weight w is kept, so that the influence of the LiDAR channel can be reduced when surface conditions are known or suspected to be poor.

3) Probabilistic Classification from Intensity: The normalised residuals from the reference networks are used to build a probabilistic model of intensity for each material. For material $m \in \mathcal{M}$, the conditional density of $\mathcal{I}_R$, evaluated at a residual value ξ, is estimated by kernel density estimation (KDE) with a Gaussian kernel and Silverman's bandwidth h used for all material classes:

$$\hat{f}_m(\xi) = \frac{1}{N_m\ h} \sum_{i=1}^{N_m} \mathcal{K}\ \left(\frac{\xi - \mathcal{I}_{R,i}^{(m)}}{h} \right), \tag{13}$$

where N_m is the number of reference residuals for material m and $\mathcal{K}$ is the standard Gaussian kernel.

$$\mathcal{K}(u) = \frac{1}{\sqrt{2\pi}} \exp\ \left(-\frac{u^2}{2} \right). \tag{14}$$

KDE is chosen over a parametric model because the residual distributions can be skewed or multimodal, as visible in Figure 2(a). Each KDE is evaluated on a grid of 800 points spanning the 1st-99th percentile range of the pooled reference residuals, extended by a fixed margin to improve tail visualisation and avoid truncation at the plot limits.

Given a target segment containing N intensity residual observations $\{\mathcal{I}_{R,n}\}_{n=1}^{N}$, the mean log-likelihood under material m is:

$$\bar{\ell}_I(m) = \frac{1}{N} \sum_{n=1}^{N} \log \hat{f}_m(\mathcal{I}_{R,n}). \tag{15}$$

Dividing by N ensures that the intensity evidence is expressed on a per-observation scale regardless of segment length, which is essential for meaningful combination with the distance term in Section III-D. Applying Bayes' rule under the uniform prior of Section III-A gives the intensity-only posterior:

$$\mathcal{P}_I(m) = \frac{\exp\big(\bar{\ell}_I(m)\big)}{\sum_{m' \in \mathcal{M}} \exp\big(\bar{\ell}_I(m')\big)}. \tag{16}$$

C. IMU Joint Detection Classification

While the LiDAR intensity model captures optical surface properties, a complementary sensing modality based on a different physical mechanism can resolve ambiguities where intensity alone is insufficient. This subsection introduces such a modality based on IMU data that exploits the mechanical response of the robot as it traverses discontinuities at pipe joints. The key observation is that different pipe materials are manufactured in characteristic segment lengths that depend on material, though the specific length also varies with manufacturer and era of installation. The nominal values used here (2.0 m for clay, 2.5 m for concrete, 3.1 m for plastic) are representative of the test networks. Measuring inter-joint distances thus provides evidence about the underlying material. This

mechanical signal is entirely insensitive to surface conditions such as moisture, fouling, and biofilm, which are the primary confounds of the intensity channel.

1) Physical Principle and Joint Detection: When the robot crosses a pipe joint, the geometric discontinuity produces a transient response in the IMU orientation signal, most prominently in the pitch axis. Let $\varphi(x)$ denote the pitch angle as a function of longitudinal position. The raw signal contains both slow drift from the pipe geometry and sharp transients from joint crossings. A moving median is subtracted to isolate the perturbations:

$$\varphi_{\mathrm{HP}}(x) = \varphi(x) - \tilde{\varphi}(x), \tag{17}$$

where $\tilde{\varphi}(x)$ is the moving median of φ evaluated over a symmetric sample window centred at x. The signal is then squared to produce a non-negative energy envelope that amplifies larger perturbations relative to noise regardless of their sign:

$$\mathcal{S}(x) = \varphi_{\mathrm{HP}}(x)^2. \tag{18}$$

Joint positions are identified as local maxima of $\mathcal{S}(x)$ satisfying two constraints: a peak height exceeding the adaptive threshold $h_{\min}$, set to the 85th percentile of non-zero signal values, and a minimum separation of $d_{\min} = 1.5\,\mathrm{m}$ to prevent double detections from signal ringing within a single joint event.

2) Probabilistic Classification from Distance: Once the joints have been detected, the inter-joint distances are denoted $\mathbf{d} = [d_1, \ldots, d_{n_d}]$, where n_d is the number of detected gaps. The simplest model assumes that each d_i corresponds to one pipe segment, but this assumption breaks down when a joint is missed: the corresponding gap then spans two or more segments and should not be discarded as noise.

To handle this, we allow each observed gap to span $k = 1, \ldots, K$ pipe segments. The case $k = 1$ corresponds to a normal gap between neighbouring joints, while $k > 1$ corresponds to a gap that contains one or more missed joints. For material m with expected segment length L_m and fixed common distance tolerance σ, the likelihood of one observed gap is:

$$p(d_i \mid m) = (1-\epsilon) \sum_{k=1}^{K} \pi(k)\, \mathcal{N}\big(d_i \mid \mu_{k,m}, \sigma_k^2\big) + \epsilon\, \mathcal{U}(d_i), \tag{19}$$

where K is the largest number of segments allowed in a single observed gap, $\pi(k)$ is the prior probability that the gap spans k pipe segments, $\mathcal{U}(d_i)$ is a uniform outlier density used for unreliable observed gaps, and ϵ is a small outlier probability. For material m, a gap spanning k pipe segments has mean $\mu_{k,m} = kL_m$ and variance $\sigma_k^2 = k\sigma^2$, where L_m is estimated from the reference run and σ is a fixed common distance tolerance informed by the reference data. The outlier term limits the influence of occasional unreliable gaps caused by false detections, boundary effects, partial segments, or odometry errors. The distance uncertainty is not re-estimated separately for each candidate material on the target run, which prevents an incorrect material from fitting the target gaps by increasing its own tolerance.

The prior probability that a detected gap spans k pipe segments, denoted $\pi(k)$, is modelled with a truncated geometric distribution:

$$\pi(k) = \frac{(1-\rho)\rho^{k-1}}{\sum_{j=1}^{K}(1-\rho)\rho^{j-1}}, \tag{20}$$

where $\rho \in [0, 1)$ controls the penalty for missed joints: most prior mass is placed on $k = 1$ while non-zero probability is kept for $k > 1$ so that missed-joint cases can still be explained. In our experiments we use $K = 3$, which covers the longest gaps observed in the test networks, and $\rho = 0.3$.

Assuming that the observed gaps are conditionally independent given the material m, the total distance log-likelihood is:

$$\ell_d(m) = \sum_{i=1}^{n_d} \log p(d_i \mid m). \tag{21}$$

As in the LiDAR model, we use the mean log-likelihood:

$$\bar{\ell}_d(m) = \frac{1}{n_d}\, \ell_d(m), \tag{22}$$

so that the score is not dominated by the number of detected gaps. The distance-only posterior is then obtained using Bayes' rule with a uniform prior:

$$\mathcal{P}_d(m) = \frac{\exp\big(\bar{\ell}_d(m)\big)}{\sum_{m' \in \mathcal{M}} \exp\big(\bar{\ell}_d(m')\big)}. \tag{23}$$

The mixture model also provides a diagnostic for each gap. For material m, conditional on the gap being explained by the pipe segment mixture, the probability that d_i spans k pipe segments is:

$$p(k \mid d_i, m) = \frac{\pi(k)\, \mathcal{N}\big(d_i \mid \mu_{k,m},\, \sigma_k^2\big)}{\sum_{j=1}^{K} \pi(j)\, \mathcal{N}\big(d_i \mid \mu_{j,m},\, \sigma_j^2\big)}. \tag{24}$$

A high probability for $k > 1$ indicates that the gap is likely to contain one or more missed joints. Thus, a long gap can still support the correct material model when it is close to an integer multiple of the expected segment length.

D. Sensor Fusion

Having defined both modalities within the shared Bayesian framework, this subsection describes how they are combined. The intensity channel contributes $N \sim 10^4$ to 10^5 residual observations per classification window while the distance channel contributes only $n_d \sim 5$ to 10 inter-joint measurements, so combining raw log-likelihood sums directly would let intensity dominate by sheer volume. The fusion therefore uses the mean log-likelihoods $\bar{\ell}_I(m)$ and $\bar{\ell}_d(m)$ defined in (15) and (22), which remove the sample-count imbalance. The intensity and distance log-densities still live on different scales, so a scalar weight w is retained to absorb this remaining scale difference. The fused log-posterior score for material m is then:

$$\ell_{\mathrm{fused}}(m) = \bar{\ell}_d(m) + w\ \bar{\ell}_I(m) + \log p(m), \tag{25}$$

where $w \geq 0$. The weighting is asymmetric by convention so that w is interpreted directly as the trust placed in the intensity channel at equal per-observation contributions. To interpret w, define the geometric-mean likelihoods:

$$\bar{\mathcal{L}}_d(m) = \exp\big(\bar{\ell}_d(m)\big), \qquad \bar{\mathcal{L}}_I(m) = \exp\big(\bar{\ell}_I(m)\big), \quad (26)$$

so that (25) corresponds to the posterior:

$$\mathcal{P}_{\text{fused}}(m) \propto \bar{\mathcal{L}}_d(m) \cdot \big[\bar{\mathcal{L}}_I(m)\big]^w \cdot p(m). \quad (27)$$

Setting $w = 1$ weights both modalities equally per observation. Values $w < 1$ reduce the influence of intensity relative to distance, which may be appropriate when surface fouling degrades the optical signal. The fused posterior is recovered from ℓ_{fused} via the log-sum-exp identity (4).

1) Target-Run Classification: For each target run, the detected joints split the run into segments. The LiDAR residuals inside each segment are used by the intensity model, and the gaps between joints are used by the distance model. The two likelihoods are then combined into a single posterior over material classes. Let $b_0 < b_1 < \cdots < b_{N+1}$ be the positions that split the run into $N+1$ segments: b_0 is the run start, b_{N+1} is the run end, and $b_1, \ldots, b_N$ are the detected joints. For each segment $[b_i, b_{i+1}]$, $i = 0, \ldots, N$, the intensity residuals are collected as:

$$\mathcal{I}_{R,n} : x_{\text{hit},n} \in [b_i, b_{i+1}], \quad (28)$$

and the mean intensity log-likelihood $\bar{\ell}_I(m)$ is evaluated by querying the reference KDE models. The mean distance log-likelihood $\bar{\ell}_d(m)$ is computed from (21) and (22), so that each observed gap is treated as a single segment ($k = 1$), a missed-joint case ($k > 1$), or an unreliable gap through the outlier component. The fused posterior (25) is then computed, and the classification decision (5) is applied.

IV. Experimental Setup

A. Pipe Networks

All experiments are conducted on three single-material pipe networks of 300 mm diameter, each consisting of consecutive segments of a single known material: concrete (segment length 2.5 m), plastic (3.1 m), and clay (2.0 m). Each network was traversed twice under the same conditions. The first run serves as the reference, from which the material-specific KDE intensity distributions and distance models are built. The second run, collected independently, serves as the classification target on which fusion is evaluated. Because reference and target runs are drawn from the same physical networks, the experiments provide a controlled test of repeatability rather than generalisation to a new sewer network. Essentially, they show whether the proposed probabilistic cues are repeatable under controlled conditions, but they do not show whether a model trained on one network transfers directly to another network with different manufacturers, ageing, fouling, or repair history. Field validation across operational sewer networks is therefore needed as a next step. Table I summarises the six datasets.

The inter-joint distance statistics extracted from each reference run are as follows. The concrete network yields $n_d = 7$ distance measurements with sample mean $\bar{d} = 2.52$ m and standard deviation $s_d = 0.13$ m. The plastic network yields $n_d = 4$ measurements ($\bar{d} = 3.24$ m, $s_d = 0.60$ m), and the clay network yields $n_d = 2$ measurements ($\bar{d} = 2.03$ m, $s_d = 0.13$ m). In this controlled dataset, the separation between the concrete and clay mean lengths is 0.49 m, while the plastic mean is 0.72 m longer than concrete. These gaps are sufficient for classification in this work, but in operational networks segment lengths can vary within a material due to in-situ repairs, motivating the fusion approach.

TABLE I
Summary of the six datasets.

Run	Material	Length	Joints	Role
Concrete 1	Concrete	20.0 m	8	Reference
Concrete 2	Concrete	20.0 m	8	Target
Plastic 1	Plastic	15.0 m	5	Reference
Plastic 2	Plastic	15.0 m	5	Target
Clay 1	Clay	6.0 m	3	Reference
Clay 2	Clay	6.0 m	3	Target

B. Robotic Platform

The inspection robot is a commercial Minicam crawler carrying several sensors, of which this work uses the LD06 LiDAR scanner mounted at a fixed 30° tilt, an IMU-BNO085 providing orientation quaternions, and wheel odometry. The LiDAR is rigidly calibrated to the robot body frame, and odometry transforms are used to associate each scan and IMU reading with the robot's longitudinal position along the pipe axis, which serves as the common spatial coordinate.

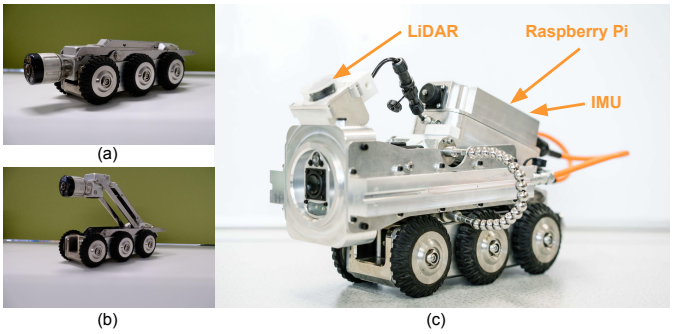


Fig. 1. The sewer inspection robot used in this study, equipped with the LD06 LiDAR scanner, IMU-BNO085, and wheel odometry.

All data are recorded as ROS 2 bags with nanosecond-precision timestamps. Temporal alignment between sensor messages and odometry transforms uses nearest-neighbour matching with a maximum tolerance of 0.5 s. Each LiDAR scan and IMU reading is associated with the robot's longitudinal position x, taken as the x-component of the odometry translation, which serves as the spatial coordinate along the pipe axis.

V. Results

This section evaluates the classification performance of each modality individually and of their fusion on the three target pipe networks. All results use the reference models built from the first traversal of each network, applied to the independently collected second traversal.

A. Intensity-Only Classification

Figure 2(a) shows the empirical distributions of $\mathcal{I}_R$ for each pipe material built from the reference networks. Clay produces a narrow, concentrated distribution shifted toward positive residuals, indicating consistently higher reflectance than the geometric baseline predicts. Concrete is broader and centred near zero. Plastic exhibits the greatest spread, with a low, wide profile that overlaps substantially with both other materials.

To quantify the discriminative power of the intensity channel, the intensity-only posterior $\mathcal{P}_I(m)$ from (16) is evaluated on each target pipe network. Figure 2 and the LiDAR column of Table II present the results.

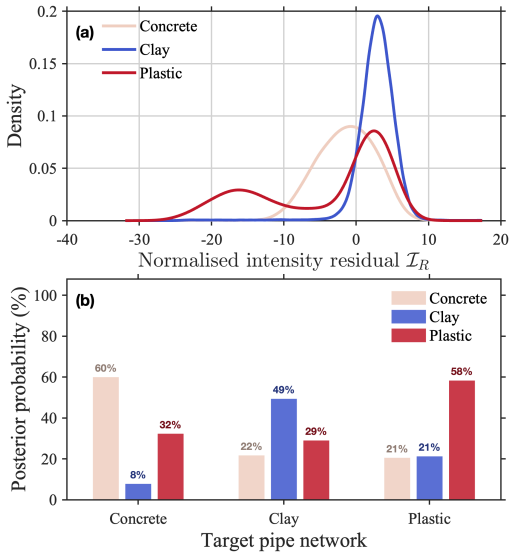


Fig. 2. Intensity-only classification using the shared polynomial residual $\mathcal{I}_R$. (a) Reference KDE for each material. (b) Posterior probabilities on the three target networks. All targets are correctly classified but with limited confidence (49-60%).

All three networks are correctly classified in the sense that the true material receives the highest posterior, but the confidence is limited: concrete is identified with 59.9% posterior probability, clay with 49.4%, and plastic with 58.3% (Table II, LiDAR column). In each case, substantial posterior mass is distributed across the competing materials. For the clay target, 29.0% of the posterior is assigned to plastic and 21.7% to concrete, producing near-maximum uncertainty for a three-class problem. These results confirm that intensity residuals carry material information but cannot support confident classification in isolation.

B. Distance-Only Classification

Figure 3 shows the joint detection results across all three reference networks. The concrete network produces eight clearly separated peaks at approximately 2.5 m spacing ($\bar{d} = 2.52$ m, $s_d = 0.13$ m, $n_d = 7$ distances). The clay network produces two strong detections consistent with the expected 2.0 m segment length ($\bar{d} = 2.03$ m, $s_d = 0.13$ m, $n_d = 2$ distances). The plastic network yields four detections with greater variability ($\bar{d} = 3.24$ m, $s_d = 0.60$ m, $n_d = 4$ distances), reflecting the wider spread of plastic joint perturbation amplitudes.

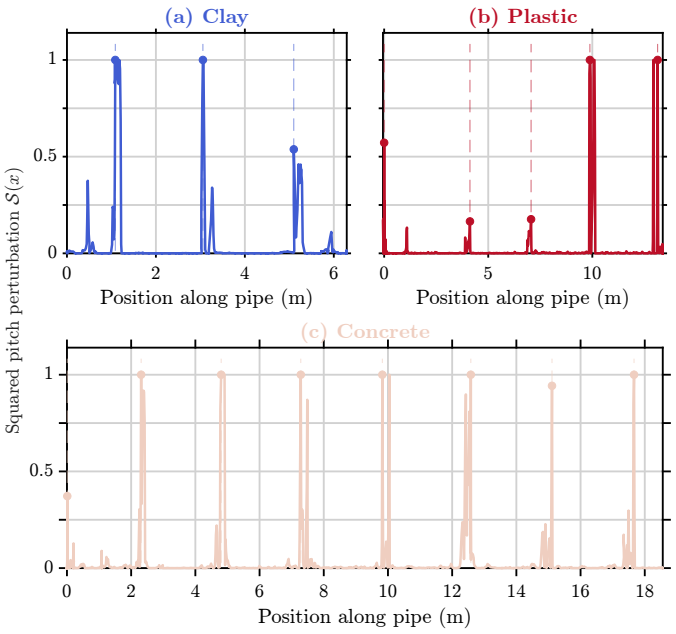


Fig. 3. Joint detection from the squared pitch perturbation signal $\mathcal{S}(x)$ for the three reference pipe networks. Filled circles mark detected joint positions. Spacing between consecutive joints is consistent with the nominal segment length of each material: 2.0 m (clay), 2.5 m (concrete), and 3.1 m (plastic).

Figure 4 visualises the nominal single-segment distance cue ($k = 1$) for the three materials, alongside the measured inter-joint distances from all reference networks. Although the figure shows only the single-segment case, classification uses the full mixture model in (19), which also accounts for multi-segment gaps caused by missed joints. The distance channel provides strong separation between all three materials, in direct contrast to the overlapping intensity distributions in Figure 2(a).

The distance channel assigns the highest posterior probability to the correct material in all cases, but with moderate confidence: 73.8% for concrete, 86.5% for clay, and 69.0% for plastic. These values are lower than in the single-Gaussian distance model because the revised likelihood is more conservative: it allows missed-joint explanations, includes an outlier component, and avoids re-estimating a separate uncertainty for each candidate material on the target run. This prevents the IMU channel from becoming overconfident when only a small number of detected gaps is available or when one gap is unreliable.

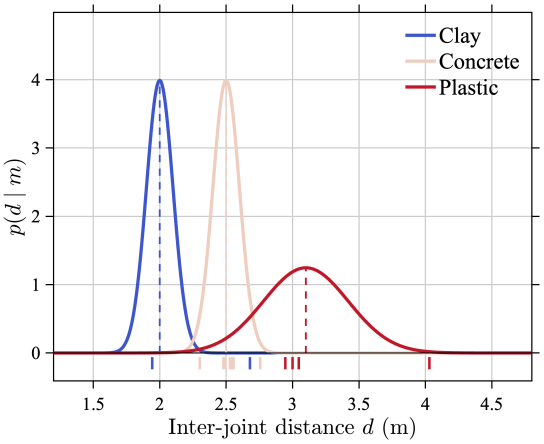


Fig. 4. Distance likelihoods under the single-segment assumption ($k = 1$) for each material. Dashed vertical lines mark the expected segment length L_m, and tick marks below the axis show the inter-joint distances measured in the reference runs. During classification, the model also evaluates the multi-segment terms $k > 1$ in (19), so a gap close to $\mu_{2,m} = 2L_m$ or $\mu_{3,m} = 3L_m$ contributes evidence as a likely missed-joint case. Gaps that are not close to a plausible multiple are instead down-weighted by the outlier component.

The mixture model in (21) also produces a per-gap diagnostic through the probabilities in (24). Most gaps in the reference runs concentrate their mass on $k = 1$, indicating that the detected joints are genuinely neighbouring. This is consistent with the clean joint detection visible in Fig. 3. When a joint is missed, a longer gap can shift probability onto $k = 2$ or $k = 3$ if it is close to an integer multiple of the expected segment length. If a gap is not close to any plausible multiple, the outlier component reduces its influence on the posterior.

C. Fused Classification

Table II presents the complete classification results on the three target runs, comparing distance-only, intensity-only, and fused conditions at $w = 1.0$. All three target networks are correctly classified under all three conditions, indicating that each modality alone is sufficient for coarse material discrimination in this controlled setting. The value of fusion lies not in changing the predicted class, but in reallocating posterior mass away from competing materials when one channel is ambiguous.

TABLE II
CLASSIFICATION RESULTS ON THE TARGET PIPE NETWORKS.

Target	Material	IMU	LiDAR	Fused
Concrete	Clay	21.9%	7.8%	3.6%
	Concrete	**73.8%**	**59.9%**	**93.5%**
	Plastic	4.3%	32.3%	2.9%
Clay	**Clay**	**86.5%**	**49.4%**	**93.5%**
	Concrete	12.9%	21.7%	6.1%
	Plastic	0.6%	29.0%	0.4%
Plastic	Clay	8.9%	21.2%	4.1%
	Concrete	22.1%	20.5%	9.7%
	Plastic	**69.0%**	**58.3%**	**86.2%**

For the concrete target, both sensing modalities favour concrete with moderate confidence: 73.8% from IMU distance and 59.9% from LiDAR intensity. The intensity channel leaks substantial probability to plastic (32.3%) and a small amount to clay (7.8%), while the distance channel spreads residual mass between clay (21.9%) and plastic (4.3%). Fusion gives a posterior of 93.5% on concrete with 3.6% assigned to clay and 2.9% to plastic.

The clay target shows the strongest complementarity. While intensity alone produces near-maximum uncertainty for a three-class problem (49.4% clay, 29.0% plastic, 21.7% concrete), the distance channel is more confident, assigning 86.5% to clay. Fusion raises the clay posterior to 93.5%, with 6.1% on concrete and 0.4% on plastic, demonstrating how the distance evidence pulls posterior mass away from the competing classes that intensity alone could not resolve.

For the plastic target, both modalities favour plastic but with moderate confidence: 69.0% from IMU and 58.3% from LiDAR. The reduced IMU confidence is caused by one observed gap of 4.12 m, which is not close to either one or two nominal plastic segment lengths. The revised distance model therefore limits the influence of this gap rather than forcing it into a missed-joint explanation. The remaining three gaps are close to the expected plastic segment length and support the plastic class. Fusion combines this agreement and raises the plastic posterior to 86.2%, with residual mass on concrete (9.7%) and clay (4.1%).

VI. Discussion and Conclusion

The results confirm that LiDAR intensity and IMU inter-joint distance provide complementary information. Their failure modes are largely orthogonal: intensity is sensitive to surface optical conditions and produces only moderate confidence, with substantial posterior mass assigned to competing materials. Distance is insensitive to surface conditions but is limited by the number and reliability of detected joints. Fusion improves robustness by combining both sources of evidence, while keeping uncertainty when one channel is less reliable, as seen in the plastic case.

The strong performance of the IMU channel in this study reflects the fact that the three tested materials have distinct nominal segment lengths in a controlled laboratory setting. In live sewer networks, segment lengths can vary within a material due to differences in manufacturer, cut pipes and in-situ repairs, reducing the power of distance alone. Under such conditions, the complementary intensity channel becomes essential and the fusion framework is designed to accommodate this by shifting posterior mass according to whichever modality provides stronger evidence.

Normalising log-likelihoods by sample count removes the imbalance between the two channels: intensity contributes $\sim 10^4 - 10^5$ residuals per segment, while distance contributes only $\sim 5 - 10$ gaps, so without normalisation intensity would dominate the fused score regardless of w. After normalisation, w instead controls the relative trust between channels at the per-observation level. The LiDAR correction only accounts for geometry, not all surface conditions. Moisture, dirt, and biofilm can still affect the intensity measurements. In such

cases, setting $w < 1$ provides a principled way to reduce the influence of the LiDAR channel and lean more on the distance channel.

Although this study uses equal prior probabilities for clay, concrete, and plastic, the same Bayesian framework could also include soft prior information from asset records, pipe age, diameter, nearby materials, or repair history. Such information must be used carefully because sewer records can be incomplete or wrong.

Several limitations bound the scope of the current results. The experimental validation uses laboratory networks with clean surfaces and single known materials per run, whereas operational sewers contain dirt, moisture, fouled surfaces, standing water, mixed materials, repairs, and irregular segment lengths. The two sensing modalities fail independently, supporting robustness, but reference and target runs traverse the same pipes, testing repeatability rather than generalisation. The number of inter-joint distance samples per reference network is small, so the estimated segment lengths L_m may be influenced by individual detections. With larger field datasets, pooled or Bayesian estimators could provide a more systematic way to estimate L_m and select the common distance tolerance σ.

Joint detection is also a boundary condition: The mixture model in (21) helps handle missed joints by allowing each observed gap to represent up to K pipe segments. Therefore, a long gap close to kL_m can still support material m instead of being treated as noise. The per-gap probabilities in (24) can also indicate likely missed joints. The outlier component reduces the influence of gaps that are not close to any plausible multiple, but false joint detections, cut pipes, or local repairs with non-standard lengths may still reduce classification confidence.

In summary, this paper presents a probabilistic framework that fuses LiDAR intensity with IMU-derived distances between detected joints for pipe material classification. A shared bivariate polynomial normalises intensity for range and scan angle effects. Residuals from the shared normalisation are used to fit one KDE per material. A distance model compares each inter-joint gap with integer multiples of the expected segment length and limits the influence of unreliable gaps through an outlier component. Both channels are combined in the log domain via mean log-likelihoods, with a tunable weight. On the three target pipe networks, a sensor data fusion algorithm assigns posterior probabilities of 93.5%, 93.5%, and 86.2% to the true concrete, clay, and plastic classes, respectively, exceeding both the LiDAR-only posteriors (59.9%, 49.4%, 58.3%) and the IMU-only posteriors (73.8%, 86.5%, 69.0%) in every case..

Future work will focus on field validation across operational sewer networks with fouling, repairs, mixed materials, and manufacturer-dependent segment lengths, where the practical value of the proposed approach will be validated.

REFERENCES

[1] S. Saegrov, Computer Aided Rehabilitation of Sewer and Storm Water Networks - CARE-S. IWA Publishing, 11 2006.

[2] Dirksen, J., Clemens, F. H. L. R., Korving, H., Cherqui, F., Le Gauffre, P., Ertl, T., Snaterse, C. T. M. (2013). The consistency of visual sewer inspection data. Structure and Infrastructure Engineering, 9(3), 214–228.

[3] Tscheikner-Gratl, F., Caradot, N., Cherqui, F., Leitão, J. P., Ahmadi, M., Langeveld, J. G., Clemens, F. (2019). Sewer asset management – state of the art and research needs. Urban Water Journal, 16(9), 662–675.

[4] C. Egger, A. Scheidegger, P. Reichert, and M. Maurer, 'Sewer deterioration modeling with condition data lacking historical records', Water Research, vol. 47, no. 17, pp. 6762–6779, 2013.

[5] A. A. F. Nassiraei, Y. Kawamura, A. Ahrary, Y. Mikuriya and K. Ishii, "A New Approach to the Sewer Pipe Inspection: Fully Autonomous Mobile Robot "KANTARO"," IECON 2006 - 32nd Annual Conference on IEEE Industrial Electronics, Paris, France, 2006, pp. 4088-4093.

[6] M. R. Halfawy, J. Hengmeechai, and T. Ormsby, "Efficient algorithm for crack detection in sewer images from closed-circuit television inspections," *J. Infrastr. Syst.*, vol. 14, no. 4, pp. 312–318, 2008.

[7] Yicheng Yu, Rob Worley, Sean Anderson, Kirill V. Horoshenkov; Microphone array analysis for simultaneous condition detection, localization, and classification in a pipe. J. Acoust. Soc. Am. 1 January 2023; 153 (1): 367–383.

[8] T. Tian et al., "Visual-Inertial-Laser-Lidar (VILL) SLAM: Real-Time Dense RGB-D Mapping for Pipe Environments," 2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), Detroit, MI, USA, 2023, pp. 1525-1531.

[9] Karnezis, Aristeidis et al. (2025). Feature Detection and Classification in Buried Pipes using LiDAR: Dataset with Camera–LiDAR Synchronisation. The University of Sheffield. Collection.

[10] Jutzi, Boris & Gross, H.. (2009). Normalization Of Lidar Intensity Data Based On Range And Surface Incidence Angle. ISPRS Ann. Photogramm. Remote Sens. Spat. Inf. Sci.. 38.

[11] B. Höfle and N. Pfeifer, 'Correction of laser scanning intensity data: Data and model-driven approaches', ISPRS Journal of Photogrammetry and Remote Sensing, vol. 62, no. 6, pp. 415–433, 2007.

[12] W. Y. Yan and A. Shaker, "Radiometric Correction and Normalization of Airborne LiDAR Intensity Data for Improving Land-Cover Classification," in IEEE Transactions on Geoscience and Remote Sensing, vol. 52, no. 12, pp. 7658-7673, Dec. 2014.

[13] Aitken, J., Evans, M., Worley, R. et al. (2021) Simultaneous localization and mapping for inspection robots in water and sewer pipe networks: a review. IEEE Access, 9. pp. 140173-140198. ISSN: 2169-3536

[14] A. C. Murtra and J. M. Mirats Tur, "IMU and cable encoder data fusion for in-pipe mobile robot localization," 2013 IEEE Conference on Technologies for Practical Robot Applications (TePRA), Woburn, MA, USA, 2013, pp. 1-6.

[15] Edwards S, Zhang R, Worley R, Mihaylova L, Aitken J and Anderson SR (2023) A robust method for approximate visual robot localization in feature-sparse sewer pipes. Front. Robot. AI 10:1150508.

[16] Sahli, Hussein and El-Sheimy, Naser, A novel method to enhance pipeline trajectory determination using pipeline junctions, Sensors (Switzerland), 2016

[17] Guan, Lianwu and Xu, Xu and Gao, Yanbin and Liu, Fanming and Rong, Hanxiao and Wang, Meng and Noureldin, Aboelmagd, "Micro-inertial-aided high-precision positioning method for small-diameter PIG navigation," in Advances in Human and Machine Navigation Systems, 2018

[18] S. Kaasalainen, A. Jaakkola, M. Kaasalainen, A. Krooks, and A. Kukko, "Analysis of incidence angle and distance effects on terrestrial laser scanner intensity," *Remote Sensing*, vol. 3, no. 10, pp. 2207–2221, 2011.

[19] A. G. Kashani, M. J. Olsen, C. E. Parrish, and N. Wilson, "A review of LiDAR radiometric processing: from ad hoc intensity correction to rigorous radiometric calibration," *Sensors*, vol. 15, no. 11, pp. 28099–28128, 2015.

[20] S. Kaasalainen et al., "Radiometric Calibration of LIDAR Intensity With Commercially Available Reference Targets," in IEEE Transactions on Geoscience and Remote Sensing, vol. 47, no. 2, pp. 588-598, Feb. 2009.

[21] Tatoglu, A. and Pochiraju, K. (2012). Point cloud segmentation with LIDAR reflection intensity behavior. In 2012 IEEE International Conference on Robotics and Automation, pp. 786–790.

[22] S. Thrun, W. Burgard, and D. Fox, *Probabilistic Robotics*. Cambridge, MA: MIT Press, 2005.

[23] C. M. Bishop, *Pattern Recognition and Machine Learning*. New York: Springer, 2006.

[24] O. Duran, K. Althoefer and L. D. Seneviratne, "State of the art in sensor technologies for sewer inspection," in IEEE Sensors Journal, vol. 2, no. 2, pp. 73-81, April 2002.