



Universiteit
Leiden
The Netherlands

Machine learning and computer vision for urban drainage inspections

Meijer, D.W.J.

Citation

Meijer, D. W. J. (2023, November 7). *Machine learning and computer vision for urban drainage inspections*. Retrieved from <https://hdl.handle.net/1887/3656056>

Version: Publisher's Version

License: [Licence agreement concerning inclusion of doctoral thesis in the Institutional Repository of the University of Leiden](#)

Downloaded from: <https://hdl.handle.net/1887/3656056>

Note: To cite this publication please use the final published version (if applicable).



STEREOVISION AND GEOMETRY RECONSTRUCTION

So far, we have looked at unsupervised and supervised machine learning methods applied to image data generated by current inspection practices in urban drainage. In this chapter, we move beyond current inspection practices, and extend image acquisition to a stereovision system, consisting of two cameras, in order to reconstruct the three-dimensional pipe geometry and recognise potential defects from that. We consolidate several techniques into RADIUS (Robust Anomaly Detection In Urban drainage with Stereovision), a framework for anomaly detection in sewer pipes.

5.1 INTRODUCTION

The RADIUS framework is designed to be compatible with the existing workflow of the trained operator, as well as be future-proof for a completely automated sewer inspection system, for which advances are being made rapidly. The framework also has a low up-front investment in terms of equipment, as it uses two cameras for data collection, and image processing steps that can be performed on a consumer-grade computer within a reasonable amount of time.

Our proposed method revolves around the technique of *computer stereovision*, which uses two or more calibrated cameras placed side by side, in order to create a sense of depth, similar to how the binocular vision in humans is used to capture the spatial configuration of one's surrounding. In that sense, the proposed method is a type of 3D ranging technique that promises to produce a faithful 3D reconstruction of the interior of a sewer pipe in the form of a 3D point cloud with associated colour information.

Since the raw output of our stereovision setup captures the pipe's surface in considerable detail (the extent of this is determined by the resolution of the cameras), it can theoretically be used to recognise various different categories of pipe defects that have a spatial nature. These include deposits, holes, fissures, intrusions and exposed granulates. Some of these may in fact be harder to correctly classify using traditional single CCTV setups, since without further spatial clues, they cannot be distinguished (for example, a fissure (*on* the surface) vs. an intruding root (*away* from the surface)).

While in terms of types of defects to be recognised, the framework is open-ended, we focus in our experiments on the recognition of two specific types of defects, namely deposits and exposed granulates. For other defect types (such as misaligned joints) the larger part of our proposed pipeline remains unchanged, but in the final stages provisions need to be made to account for the different spatial phenomena at hand. When focusing on deposits and exposed granulates, we need to recognise areas in the pipe where the surface is further inward or outward than one would expect. In other words, there is an *expected* pipe geometry (the pipe model) and a *measured* pipe geometry, and any deviations between these will be expressed in an anomaly score. Although this form of anomaly detection sounds fairly straightforward, it relies on the availability of a pipe model, which is actually non-trivial. First of all, the precise location and orientation of the camera pair inside the pipe is uncertain. We can only assume that the cameras are pointing roughly along the main axis of the pipe, not too far removed from the centre of the pipe. Second, our method should be robust with respect to the shape of the pipe, such that different pipe topologies can be dealt with, without having to reconfigure the recognition system. This means we will take a data-driven approach that assumes that the larger part of the pipe is unaffected, such that a 'normal' geometry can be derived from that, and the




Figure 5.1: Overview of the proposed framework

outliers with respect to this geometry are the anomalies.

In broad strokes, our framework works as follows (see also figure 5.1). In the image acquisition stage, any radial lens distortion is removed from the images, and misalignment between the left and right image are corrected. An existing algorithm for stereo matching ¹ then produces pairs of corresponding pixels in both images. The computed disparity between each pair can be translated into a distance for this pixel: points closer to the cameras will appear further apart in the two images. The resulting point cloud already captures the pipe geometry, but needs to be further processed in order to automatically identify the defects of interest, in our case deposits and exposed granulate. The next stage of surface fitting combines a parameterised surface model, one that can adapt to a wide range of pipe types, with the robust regression algorithm RANSAC ². This algorithm assumes that the majority of the data fits a predefined model class (our parameterised pipe model), but also allows for a fraction of the data to constitute outliers. This allows us to use the parameterised pipe model without the fit being influenced by these outliers. The deviation of the measured geometry from the expected geometry as predicted by the RANSAC model is used as an anomaly score.

The main contributions of this chapter are:

- ◇ We demonstrate how a faithful and high-resolution reconstruction of the pipe surface, including its defects, can be obtained with stereo cameras and a stereo matching algorithm.
- ◇ We propose a generic pipe surface model that is able to model the pipe geometry of a range of pipe shapes (including circular and egg-shaped), captured under various angles. This pipe surface model has the attractive property that it falls in the category of functions that can be statistically fit with the Ordinary Least Squares method ³, making it computationally efficient.

¹ HIRSCHMÜLLER, H. 2005. Accurate and efficient stereo processing by semi-global matching and mutual information. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05)*. Vol. 2. IEEE, 807–814

² FISCHLER, M. A. AND BOLLES, R. C. 1981. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM* 24, 6, 381–395

³ GOLDBERGER, A. S. 1964. *Classical Linear Regression, Econometric Theory*. New York: John Wiley & Sons

- ◇ We propose a method based on the RANSAC algorithm to fit point cloud data that is a mixture of regular pipe surface and anomalies.
- ◇ We define a global anomaly score that quantifies the amount of deviation from the pipe model per image pair.

Section 5.2 outlines relevant prior work. The framework itself is described in full detail in section 5.3. Section 5.4 gives an overview of the data and the experiments, the results of which are summarised and discussed in section 5.5. Section 5.6 discusses the limitations, envisioned applications, and possible future work.

5.2 PRIOR WORK AND MOTIVATION

In the field of 3D ranging techniques, where our approach belongs, the use of laser scanners for sewer pipe inspections has been thoroughly researched, see for example ^{4,5,6,7}. The reasons we have opted to go with stereovision instead of laser scanning are threefold, i) the equipment cost of two cameras versus that of a laser scanner is significantly lower, making this approach more accessible, ii) a stereovision setup has no moving parts, which matters in real-world scenarios, where the environment of a sewer pipe can be very abrasive to moving parts in particular, and iii) the point cloud obtained from stereovision will be linked directly to images with a colour component, whereas a laser scanner only provides the geometry.

While not much research has been done on the use of stereovision in the context of sewer condition assessment, the use of stereovision in the general context of sewer maintenance is not new. Most works restrict their approach to

⁴ LEPOT, M., STANIĆ, N., AND CLEMENS, F. H. 2017. A technology for sewer pipe inspection (part 2): Experimental assessment of a new laser profiler for sewer defect detection and quantification. *Automation in Construction* 73, 1–11

⁵ STANIĆ, N., CLEMENS, F. H., AND LANGEVELD, J. G. 2017. Estimation of hydraulic roughness of concrete sewer pipes by laser scanning. *Journal of Hydraulic Engineering* 143, 2, 04016079

⁶ DURAN, O., ALTHOEFER, K., AND SENEVIRATNE, L. D. 2007. Automated pipe defect detection and categorization using camera/laser-based profiler and artificial neural network. *IEEE Transactions on Automation Science and Engineering* 4, 1, 118–126

⁷ BAHNSEN, C. H., JOHANSEN, A. S., PHILIPSEN, M. P., HENRIKSEN, J. W., NASROLLAHI, K., AND MOESLUND, T. B. 2021. 3d sensors for sewer inspection: A quantitative review and analysis. *Sensors* 21, 7, 2553

⁸ AHRARY, A., TIAN, L., KAMATA, S.-I., AND ISHIKAWA, M. 2005. An autonomous sewer robots navigation based on stereo camera information. In *17th IEEE International Conference on Tools with Artificial Intelligence (IC-TAI'05)*. IEEE, 6–pp

⁹ AHRARY, A. AND ISHIKAWA, M. 2008. A fast stereo matching algorithm for sewer inspection robots. *IEEE Transactions on electrical and electronic engineering* 3, 4, 441–448

¹⁰ KOODTALANG, W., SANGSUWAN, T., AND NOPPAKAO, B. 2018. A design of automated inspections of both shape and height simultaneously based on stereo vision and plc. In *18th International Conference on Control, Automation and Systems (ICCAS)*. IEEE, 1290–1294

¹¹ GUNATILAKE, A., PIYATHILAKA, L., KODAGODA, S., BARCLAY, S., AND VITANAGE, D. 2019. Real-time 3d profiling with rgb-d mapping in pipelines using stereo camera vision and structured ir laser ring. In *14th IEEE Conference on Industrial Electronics and Applications (ICIEA)*. IEEE, 916–921

¹² HUYNH, P., ROSS, R., MARTCHENKO, A., AND DEVLIN, J. 2015. Anomaly inspection in sewer pipes using stereo vision. In *IEEE International Conference on Signal and Image Processing Applications (ICSIPA)*. IEEE, 60–64

¹³ HUYNH, P., ROSS, R., MARTCHENKO, A., AND DEVLIN, J. 2016. 3d anomaly inspection system for sewer pipes using stereo vision and novel image processing. In *IEEE 11th Conference on Industrial Electronics and Applications (ICIEA)*. IEEE, 988–993

¹⁴ MEIJER, D. W., KESTELOO, M., AND KNOBBE, A. J. 2018. Unsupervised anomaly detection in sewer images with a PCA-based framework. In *International Conference on Pattern Recognition and Artificial Intelligence (ICPRAI)*. 354–359

cylindrical pipes, as these are fairly common, but our work does not impose that limitation.

Ahrary et al. (2005)⁸ propose an algorithm for navigation of an autonomous vehicle through a sewer network based on stereovision. Later, Ahrary et al. (2008) developed a computationally efficient stereo matching algorithm specifically for sewer pipes⁹. We do not use either of these algorithms, as processing power is not a limitation in our research, and navigation of an autonomous vehicle is outside our scope.

Tangentially related to this work, Koodtalang et al.¹⁰ use stereovision to determine manufacturing defects in pipes prior to installation.

Gunatilake et al.¹¹ combine a stereovision setup with two laser profilers to map the images recorded by the cameras onto the potentially more accurate point cloud produced by the laser profilers. This produces a high-resolution RGB-D dataset for later inspection, either by a trained expert or another algorithm. The method is tested on a single, heavily corroded pipe, as well as an artificial pipe.

Most closely related to our work, Huynh et al. have published two works^{12,13} on anomaly detection in sewer pipes with stereovision. They demonstrate the visibility of artificial defects in point clouds generated from stereovision. A critical difference to their approach is that our approach requires no human operator to centre the defect into the camera's field of view, but instead is able to highlight anomalies in the entire pipe from a set of images, and is thus more suitable for automated defect detection.

Anomaly detection has been used extensively in sewer condition assessment as a stepping stone from “traditional” data that is gathered for manual classification, towards automation of the inspection process. Meijer et al.¹⁴ performed principal component analysis (PCA) on various feature descriptors of a labeled set of CCTV images and compared the partial reconstruction with the actual val-

ues for an unsupervised approach, and compared this with a convolutional autoencoder. Myrans et al.¹⁵ performed anomaly detection on sewer CCTV images by training a random forest and a support vector machine on GIST-features. Myrans et al.¹⁶ later expanded on this by exploring the use of a one-class support vector machine, a type of support vector machine designed specifically for anomaly detection. Moradi et al.¹⁷ similarly used a one-class support vector machine to detect anomalies from SIFT features, and combined this approach with localization of the pipe through text recognition. Fang et al.¹⁸ performed anomaly detection on sewer CCTV video footage by performing principal component analysis on various local feature descriptors. Russo et al.¹⁹ use a convolutional autoencoder to detect anomalies in CCTV images.

While numerous other works that use computer vision or image processing to detect defects in sewer pipe images exist, we have limited this section to only those that perform *unsupervised anomaly detection*, as they are most similar to this work. For a more broad perspective on recent advances in this field, please refer to Haurum and Moeslund²⁰.

5.3 FRAMEWORK

We propose a framework for anomaly detection from stereovision measurements in sewer pipes, as shown in figure 5.1. The framework consists of five major steps, designed to be executed in sequence, each step's output being the next step's input. The five steps in the framework are each discussed in detail in sections 5.3.1-5.3.5.

5.3.1 IMAGE ACQUISITION

The stereovision setup requires two cameras placed side-by-side, at equal height, pointed in the same direction. Perfect

¹⁵ MYRANS, J., EVERSON, R., AND KAPELAN, Z. 2018. Automated detection of faults in sewers using cctv image sequences. *Automation in Construction* 95, 64–71

¹⁶ MYRANS, J., KAPELAN, Z., AND EVERSON, R. 2018b. Using automatic anomaly detection to identify faults in sewers. In *WDSA/CCWI Joint Conference Proceedings*. Vol. 1

¹⁷ MORADI, S., ZAYED, T., NASIRI, F., AND GOLKHO, F. 2020. Automated anomaly detection and localization in sewer inspection videos using proportional data modeling and deep learning-based text recognition. *Journal of Infrastructure Systems* 26, 3, 04020018

¹⁸ FANG, X., GUO, W., LI, Q., ZHU, J., CHEN, Z., YU, J., ZHOU, B., AND YANG, H. 2020. Sewer pipeline fault identification using anomaly detection algorithms on video sequences. *IEEE Access* 8, 39574–39586

¹⁹ RUSSO, S., DISCH, A., BLUMENSAAT, F., AND VILLEZ, K. 2020. Anomaly detection using deep autoencoders for in-situ wastewater systems monitoring data. *arXiv preprint arXiv:2002.03843*

²⁰ HAURUM, J. B. AND MOESLUND, T. B. 2020. A survey on image-based automation of cctv and sset sewer inspections. *Automation in Construction* 111, 103061

alignment of the cameras is near impossible, but correcting a slight misalignment is possible. The setup is then directed into the pipe, such that the camera axes are mostly parallel to the pipe axis. For in-situ inspection, this means that the setup has to be attached to the pipe inspection vehicle, while aimed directly into the pipe.

Any optical lens introduces some radial distortion to an image ²¹, meaning that points at different distances from the lens axis have different levels of magnification. As the lens axes of the two cameras are parallel but translated, this may introduce a difference in magnification of a point between the two cameras, and thereby a difference in vertical position. Depending on the severity of this distortion (or if the cameras and lenses are not of identical make), correction may be required for the images to be suitable for stereo matching.

By taking pictures of a chessboard pattern from different angles and distances, we can observe the effects of the radial distortion: without any distortion, the lines on a chessboard should be entirely straight, but slight curves may appear as a result of the radial distortion. Radial distortion can be reversed digitally by performing a second radial distortion to undo the first. The correct inverse distortion parameters can be estimated from the deviations in the images of the chessboard pattern, as outlined in more detail in ²².

Once images from both cameras are free of (extreme) lens distortion, an alignment between the cameras must also take place, in order to compensate for a vertical misalignment or rotation of one camera around its axis. From a set of images with several visible landmarks (the same images of the chessboard pattern may be used), we can estimate any vertical shift between the images, as well as a possible rotation along the camera axis, and correct this with a simple affine transformation ²³. If the camera axes themselves are not perfectly aligned, this may be visible as a horizontal shift of points that are very far away, as a point on the horizon

²¹ HECHT, E. ET AL. 2002. *Optics*. Vol. 5. Addison Wesley San Francisco

²² WU, Y., JIANG, S., XU, Z., ZHU, S., AND CAO, D. 2015. Lens distortion correction based on one chessboard pattern image. *Frontiers of Optoelectronics* 8, 3, 319–328

²³ BROWN, L. G. 1992. A survey of image registration techniques. *ACM computing surveys (CSUR)* 24, 4, 325–376

should in theory have the same position in both images.

After camera alignment, the first step is complete and the images can serve as input for the second step: semi-global stereo matching.

5.3.2 STEREO MATCHING

As described in section 2.4, stereo matching relies on comparing projected locations of a three-dimensional point onto two-dimensional images. Stereo matching an image is generally done by comparing positions in the reference image to horizontally shifted positions in the second image, often to sub-pixel accuracy²⁴. The shift that best matches a pixel in the reference image to a pixel in the second image is called the *disparity* for that pixel. The comparison may be done by minimising the difference in values of the pixels, but better results may be obtained by using a matching cost that relies less on absolute values, such as cross-correlation, Hamming distance, or Birchfield-Tomasi dissimilarity²⁵.

Matching single pixels from the two images is going to lead to a substantial amount of incorrect matches. Suggested solutions for this include matching a window around each pixel to a window of the same size, enforcing some type of smoothness between disparity in neighbouring pixels, and various combinations thereof²⁶.

Unique to our problem is the fact that the sewer pipe axis is parallel to the camera axes. The surface we are most interested in, the pipe wall, is perpendicular to the image plane. This causes the *Z*-distance and disparity to gradually change and not be constant anywhere inside the pipe. Window-based stereo matching methods are designed to perform best when large patches of the reference image have the same disparity, which is not the case in this scenario.

This gradual change requires the window around the pixel that is to be matched to be small (we suggest < 10 pixels on either side). The larger the window is, the more

²⁴ HIRSCHMÜLLER, H. AND GEHRIG, S. 2009. Stereo matching in the presence of sub-pixel calibration errors. In *IEEE Conference on Computer Vision and Pattern Recognition*. IEEE, 437–444

²⁵ BIRCHFIELD, S. AND TOMASI, C. 1998. A pixel dissimilarity measure that is insensitive to image sampling. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 20, 4, 401–406

²⁶ LAZAROS, N., SIRAKOULIS, G. C., AND GASTERATOS, A. 2008. Review of stereo vision algorithms: from software to hardware. *International Journal of Optomechatronics* 2, 4, 435–462

difficult it will be to match it properly. The extremes of the window are expected to have a different disparity from the centre pixel, so too large a window will be impossible to match correctly.

²⁷ HIRSCHMÜLLER, H. 2005. Accurate and efficient stereo processing by semi-global matching and mutual information. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05)*. Vol. 2. IEEE, 807–814

To enforce smoothness, we suggest using Hirschmüller's *semi-global matching* algorithm ²⁷, which adds two regularisation parameters. These parameters, P_1 and P_2 , penalise a pixel having a different disparity from its neighbouring pixels. The best match for each pixel in the reference image is chosen as the match that minimises the sum of the matching cost M and the regularisation cost R . This introduces a circular dependency, as the regularisation cost depends on the best match of neighbouring pixels, which itself depends on the regularisation cost. Because of this, the algorithm usually requires multiple passes to stabilise.

Each neighbouring pixel contributes 0, P_1 , or P_2 to the regularisation cost, depending on whether the neighbour has the same disparity, an absolute difference in disparity of 1, or a larger difference in disparity with the current pixel, respectively. Specifically, each match is given the regularisation cost:

$$R(d_1, d_2) = \begin{cases} 0 & \text{if } d_1 = d_2 \\ P_1 & \text{if } |d_1 - d_2| \leq 1 \\ P_2 & \text{if } |d_1 - d_2| > 1 \end{cases} \quad (5.1)$$

Again taking into account the fact that we expect the disparity to change gradually, we suggest setting $P_1 \ll P_2$. A value of $P_1 = 0$ may prove successful, but could also lead to more erratic results. With a small or no penalty on small differences in disparity, but a large penalty on larger disparities, we can enforce the type of smoothness that we expect in sewer pipe images. The best value of P_2 depends on the window size chosen for matching, the matching cost itself, and the number of neighbours that the semi-global matching algorithm considers (commonly 4 or 8).

To reduce false detections further, we also suggest using

a *uniqueness ratio*, which requires that the best disparity must have a score that is at least u times as large as the next best disparity. This may lead to correct disparities also being discarded, but this happens most commonly in very smooth regions, where the exact disparity is difficult to pinpoint anyway. For the purposes of anomaly detection, this is not an issue, as these regions are unlikely to contain anomalies.

With an accurate estimation of the disparity for each pixel, these disparities can be used to reconstruct a three-dimensional point cloud.

5.3.3 THREE-DIMENSIONAL GEOMETRY RECONSTRUCTION

As outlined in section 2.4, we can triangulate the three-dimensional location of a point visible in both cameras once we know the disparity, using equations (2.26), (2.27), and (2.28). Equations (2.27) and (2.28) may be rewritten as

$$X = (x - x_0) \cdot Z/f \quad (5.2)$$

$$Y = (y - y_0) \cdot Z/f \quad (5.3)$$

Where (x, y) is the pixel position of the triangulated point in the reference image, and (x_0, y_0) is the pixel position of the centre of the reference image. This means that the centre of the image will be projected to some position on the Z -axis, both the X and Y coordinates of the point being zero.

Doing this for every pixel in the image (that has a valid disparity) gives us a point cloud with one point for every pixel. It may be useful to keep the RGB values of the pixels attached to the points in the point cloud for easier inspection and later processing. But a few important caveats arise when we move away from the ideal purely mathematical situation as introduced in section 2.4 though.

The measurement of baseline b will have some non-zero error, which leads to a scaling of the entire point cloud by a factor of $\frac{b'}{b}$ where b' is the measured baseline and b the actual baseline. The more accurately the baseline is measured, the closer to 1 this scaling factor is. It should be noted that if correct physical dimensions of the point cloud are not important to the application, this does not have to be taken into consideration.

If the camera axes are not entirely parallel in the epipolar plane, the calculated disparity will have a small error. This may lead to a perceived deviation in radius along the length of the pipe, meaning a cylindrical pipe may appear conical in the point cloud.

While these are both issues to be aware of, they do not hinder the pipe model we propose in this work.

5.3.4 ROBUST PIPE SURFACE FITTING

At this point, we have a three-dimensional point cloud of a pipe which can be used to estimate the original pipe geometry as a mathematical model, excluding any anomalies (henceforth simply referred to as the ‘geometry’).

While the image will be perfectly in-focus at a specific distance, a region known as the *depth of field* around this distance is determined to be the range with *acceptable* levels of focus. The distance range of the depth of field is a simple function of the focus distance, the focal length, and the aperture size, which is adjustable in most cases. A smaller aperture size will give a larger depth of field, at the cost of less light reaching the camera sensor, leading to more sensor noise at equal exposure times²⁸. To inspect an entire pipe, we might move the pipe inspection vehicle through the pipe at small intervals, taking measurements at each interval. This means that a larger depth of field leads to fewer measurements needed to process a unit length of pipe, as a larger portion of the pipe can be captured in a single

²⁸ SALVAGGIO, N. 2009. *Basic photographic materials and processes*. Taylor & Francis

photograph. This results in a point cloud of a pipe, centred approximately around the Z -axis. Points with a Z value outside the depth of field can be discarded, as we are better off estimating the geometry of such points when the cameras are positioned at a different position along the pipe.

A transformation to cylindrical coordinates at this point allows for a more natural notation of the geometry. We define:

$$r = \sqrt{X^2 + Y^2} \quad (5.4)$$

$$\phi = \arctan_2(Y, X) \quad (5.5)$$

where $\arctan_2$ is the two-argument arctangent, which spans the interval $(-\pi, \pi]$. We can now without loss of information express each point in (r, ϕ, Z) coordinates.

A naive approach at this point might be to fit a cylinder model of

$$r = r_0 \quad (5.6)$$

to capture the geometry of the pipe, where r_0 is the inner radius of the pipe. There are a few reasons why this is a poor approach:

- i. The Z -axis may not be the precise centre of the pipe, depending on how accurately it was possible to align the reference camera's axis with the pipe axis.
- ii. The pipe might not have a circular profile. In our experiments we use both cylindrical and egg-shaped pipes, but any pipe with a somewhat smooth profile should work with our approach.
- iii. The radius and centre of the pipe may appear slanted in the point cloud along the Z -axis as a result of a slight misalignment of the camera axes in the epipolar plane.

We can address each of these issues in order.

To address the first issue, we assume for now that the pipe is cylindrical along the Z -axis, but not perfectly centred. Using a polar coordinate representation of an off-centre circle, we may express the geometry as

$$r = \sqrt{r_0^2 - d^2 \sin^2(\phi - \rho)} + d \cos(\phi - \rho) \quad (5.7)$$

where d is the distance between the axis of the pipe and the Z -axis, and ρ is the angle at which the distance to the Z -axis is maximal. It can be observed that if $d \ll r_0$ (the centre of the pipe is close to the centre of the image), we may simplify equation (5.7) to

$$r \approx r_0 + d \cos(\phi - \rho) \quad (\text{for } d \ll r_0) \quad (5.8)$$

At this point, we take a small sidestep to rewrite equation (5.8) using a trigonometric identity into

$$r = r_0 + a \sin(\phi) + b \cos(\phi) \quad (5.9)$$

It can be seen that these two forms are identical when

$$d = \sqrt{a^2 + b^2} \quad (5.10)$$

$$\rho = \arctan_2(b, a) \quad (5.11)$$

The reason for this rewrite is entirely practical: we still have two unknowns to solve for, but both unknowns are now parameters of a linear function, meaning that we can now solve a and b with Ordinary Least Squares regression ²⁹, whereas that would not be possible for the parameter ρ , because it is inside a cosine in equation (5.8).

To address the second issue, the possibility of pipes with non-circular profiles, we need a more complex function to describe the radius r as a function of the angle ϕ . To prevent modeling possible defects into the geometry, making them impossible to detect as anomalies, we use a limited approximation of r in terms of functions of ϕ . Different approximations can be used, but we suggest the use of a

²⁹ GOLDBERGER, A. S. 1964. *Classical Linear Regression, Econometric Theory*. New York: John Wiley & Sons




Figure 5.2: Cross-sectional profile of a (fictional) deformed ‘egg’-shaped pipe (black, solid), with a fit of equation (5.12) for $K = 6$ (red, dashed).

Fourier series approximation, as the radius is inherently periodic as a function of the angle.³⁰ We redefine the model as

³⁰ Angles are periodic by definition:

$$\phi \equiv \phi + 2\pi$$

$$r = r_0 + \sum_{k=1}^K \left(a_k \sin(k\phi) + b_k \cos(k\phi) \right) \quad (5.12)$$

where K dictates how many Fourier components are used to approximate the radius. It may be seen that equation (5.9) is an instance of equation (5.12), with $K = 1$. For egg-shaped pipes, we find a value of $K = 6$ to be generally sufficient, but any pipe profile with corners or otherwise non-smooth sections may require a higher value of K . Figure 5.2 illustrates how an egg-shaped pipe may be expressed in these Fourier components.

To address the third and final issue, we allow the radial parameters to change along the Z -axis, that is, along the pipe

axis. To account for both a translation and scaling of the profile along the Z -axis—corresponding to a misalignment of camera axis and pipe axis, and a measurement error in the baseline distance, respectively—we allow each of the previously introduced parameters to vary linearly along the Z -axis. For every term in equation (5.12), we add another term with a different parameter, and multiply by Z , giving us:

$$r = (r_0 + \rho_0 Z) + \sum_{k=1}^K \left(a_k \sin(k\phi) + b_k \cos(k\phi) + \alpha_k Z \sin(k\phi) + \beta_k Z \cos(k\phi) \right) \quad (5.13)$$

Equation (5.13) is the model we will use to fit the transformed point cloud data, but as we expect anomalies, we will have to employ a robust regression method. Note that, somewhat surprisingly, this model is linear in terms of the parameters to be fit, such that Ordinary Least Squares (multiple) regression can be applied.

The robust regression method we use is RANSAC, short for ‘random sample consensus’³¹. RANSAC fits a model a large number of times on a ‘minimal subset’ of data, then selects a model fit that has both a large amounts of inliers, and a low error rate for those inliers. In this context, a minimal subset is the minimum number of points we need to fit the model. As our model has $2 + 4K$ parameters, we need as many points.

Then we determine the inliers, the points that are accurately described by this fit, according to some maximum difference between the actual value of r and the one predicted by the fit, known as the *inlier threshold*. If the number of inliers meets a set minimum, the model is fit a second time, but on all inliers this time. We store this new fit, along with the error rate on its inliers. This process is repeated a large number of times, after which we select the fit with the lowest error rate on its inliers.

The value of the inlier threshold will depend on the vari-

³¹ FISCHLER, M. A. AND BOLLES, R. C. 1981. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM* 24, 6, 381–395

ance of the predicted variable. The minimum number of inliers required will usually be defined as a percentage of all data points, depending on the ratio of anomalous data points we expect to have. The chance that a minimal subset will result in a good fit of the data is low, but this first fit is only used to determine which points are the inliers that we want to perform the second fit on. Depending on how likely it is that the first fit reaches the minimum number of inliers under the chosen inlier threshold, the number of times the process should be repeated can differ by orders of magnitude: 10, 100, or 1000 could all be reasonable numbers.

If we choose the RANSAC algorithm parameters reasonably, this should give us a fit of the model described in equation (5.13) that accurately describes a large portion of the data points, while not taking the actual anomalies into account.

5.3.5 ANOMALY DETECTION AND PROCESSING

The final step in the framework, anomaly detection, is trivial at this stage: the (absolute) difference between the actual value of r and the value predicted by the best fit is an ‘anomaly score’. We might threshold these scores to distinguish anomalies from non-anomalies, or consider the scores themselves as a continuous indicator.

Depending on the context in which the framework is employed, we have several suggestions for further processing of the anomaly scores:

- ◇ the ratio of anomalous pixels to regular pixels may be an aggregated indicator of anomalousness of a length of pipe,
- ◇ for further human inspection, the anomaly scores can

be visualised in either an interactive point cloud or the original images, to indicate areas that might warrant attention,

- ◇ if pixel-wise classification is the goal, the anomaly scores can augment the RGB values of a pixel in a subsequent classifier.
- ◇ the size of connected anomalous regions, as well as the absolute anomaly scores in such regions, might be used for defect identification or even severity.

In the experiments performed for this paper, we have calculated a global anomaly score per image set as follows:

$$A = \frac{1}{N} \sum_{i=1}^N \|\min(r_i - \hat{r}_i, 0)\| \quad (5.14)$$

where r_i is the radius of a point in the point cloud and $\hat{r}_i$ is the predicted radius for that point. The point anomaly scores are clipped into the range $(-\infty, 0]$ and we calculate the average absolute value over all points in the point cloud. The point anomaly scores are clipped to only negative values, as otherwise the many points outside of the inner pipe wall present in our point clouds would skew the global anomaly score. This clipping would not be necessary in *in-situ* inspections, because no points outside the pipe should be visible, except if caused by defects.

5.4 EXPERIMENTAL SETUP

We evaluated the efficacy of the framework in a laboratory experiment. A total of 26 sewer pipe segments in various conditions were photographed with a stereocamera setup. Two Basler Ace2 A1920-160umBAS area scan cameras were fitted with lenses with a 16 millimeters focal length and attached side by side to a metal plate. The baseline was determined to be 29 millimeters, the lenses were focused at

approximately 1.5 meters distance, and the lens aperture was set to an f -number of 6, meaning the aperture diameter was equal to $16/6 \approx 2.667$ millimeters. A single pixel of an object at the in-focus distance corresponds to approximately 2.2 millimeters in real-world coordinates. In the ideal circumstances the algorithm can detect a shift of 1/16th of a pixel, so the maximum sensitivity we can expect to achieve is in the order of 1/10th of a millimeter.

The plate was attached to a rail, to allow for movement of the setup along the camera axis. The cameras were directed into the sewer pipe, which was covered at both ends with a piece of cloth. The pipe segments were illuminated with an LED light placed behind the cameras. The entire setup is depicted in figure 5.4.

22 of the 26 sewer pipes were photographed from both ends, the other 4 sewer pipes were photographed from one end only, giving us a total of 48 image sets.

5.4.1 IMAGE DATA

Figure 5.4.1 shows examples of stereo images sets of the sewer pipes as obtained with the experimental setup. Subfigure 5.4(a) shows a typical naturally aged cylindrical pipe, containing plenty of texture for the stereo matching. Subfigure 5.4(b) shows a typical naturally aged egg-shaped pipe, the reason we need a non-cylindrical model. Subfigure 5.4(c) shows a new cylindrical pipe, lacking sufficient texture to accurately stereo match.

5.4.2 IMPLEMENTATION DETAILS AND PARAMETERS

The framework was implemented in OpenCV 4.5³² and Python 3.6³³. Using the default implementation of semi-global stereo matching in OpenCV, we chose a search range

³² BRADSKI, G. 2000. The OpenCV Library. *Dr. Dobb's Journal of Software Tools*

³³ VAN ROSSUM, G. AND DRAKE, F. L. 2009. *Python 3 Reference Manual*. CreateSpace, Scotts Valley, CA




Figure 5.3: Experimental setup (not the final lighting setup).


between 20 and 220 pixels of disparity, used a block size of 7, set regularization parameters $P_1 = 100$ and $P_2 = 10,000$, and used a uniqueness ratio of $u = 10$.

After stereo matching and geometry reconstruction, the red cloth background of the images is removed with a flood fill and a valid Z -range can be selected by the user, or a default range of $Z \in [1.5, 2.0]$ meters from the camera may be used.


After conversion to cylindrical coordinates, we fit the model in equation (5.13) with $K = 6$. RANSAC is run for 10 iterations, the initial fit is calculated on 50 randomly selected datapoints, inliers are determined by a maximum absolute difference of 0.005, and the second fit is calculated on the inliers if those inliers make up at least 90% of the datapoints. The fit with the lowest error on its inliers is



(a) A typical naturally aged cylindrical sewer pipe.



(b) A typical naturally aged egg-shaped pipe.



(c) A typical new cylindrical sewer pipe.

Figure 5.4: Three examples of stereo image sets as obtained with the experimental setup. All images have lens distortion correction, the left images also have a vertical translation correction according to the process described in section 5.3.1.

chosen, or if no initial fit had enough inliers, the entire process is repeated with 10 iterations.

The best fit from the RANSAC model is applied to all data points, including those outside the Z -range that the model was fit on, and the deviation from the fit is presented for visual inspection in both a point cloud and the reference image.

The code of our implementation is available to try as a demo. It can be found at:

<https://github.com/data-flux/StereoDemo>

5.5 RESULTS AND DISCUSSION

5.5.1 STEREO MATCHING AND GEOMETRY RECONSTRUCTION

We start our discussion of the results by considering the first half of our approach, the stereo matching and geometry reconstruction. Getting an objective, unequivocal assessment of the produced point cloud is challenging, since we do not have a golden standard measurement of the 3D pipe geometry to compare the point clouds with. Because of this, we will mostly have to rely on subjective validation of the results. We asked a human assessor to judge each point cloud on how accurate and consistent the point cloud reconstructed the original 3D geometry, on a scale from 1 (very poor) to 10 (very accurate). Over the 48 image sets, an average rating of 8.4 was assigned, indicating that the 3D reconstruction was quite good. In figure 5.5.1, four examples of different types of pipes are given where the reconstruction was successful (average rating of 8.5). The left picture shows the left image of each stereo pair, and the right image shows the point cloud with points coloured by the colour from the original (left) image. The characteristics of the virtual

lens of the generated image deviate from the physical lens somewhat, and the virtual camera was deliberately placed somewhat further ahead. This allows the viewer to somewhat appreciate the 3D nature of the point cloud (rather than reproducing the images on the left without knowing the depth). Specifically, the forward camera position allows observing any occluded areas by ‘seeing around’ the humps on the pipe surface. Especially in the point cloud image subfigure 5.5.1(a), areas of occlusion behind the various deposits are clear. Viewed from the side, this point cloud has considerable gaps behind all deposits.

A minority of the pipes were not reconstructed correctly. Five image pairs (from four pipes) scored a value below 8 (average score 6), with the lowest being two scores of 5. What these five image pairs have in common, compared to the remaining successful image sets, is that they involve pipes with areas of smooth and monochrome surface, as can be seen in figure 5.5.1. Such areas are especially common in relatively new pipes which do not show much deterioration. The surface of such pipes will be hard to stereo-match, since few surface features can be used to match pixels in the stereo pair. The result is that entire patches of pixels have an undetermined depth. Further contributing factors to this poor matching are low lighting and lack of lens focus. With improved focus and lighting, the tiniest surface features may lead to proper matching again.

5.5.2 SURFACE FITTING AND ANOMALY DETECTION


Next, we consider the quality of the surface fitting and subsequent anomaly detection steps of our framework on the 48 image sets. In figure 5.5.3, the four pipes of figure 5.5.1 are shown again, with the local anomaly score assigned to the derived point cloud, colourcoded as blue being an anomaly



(a)



(b)




(c)



(d)

Figure 5.5: Four examples of the stereo matching. Left shows the reference image, right shows the reconstructed point cloud.




(a)



(b)



(c)



(d)

Figure 5.6: Images of pipes that showed problems with stereo matching, due to patches of pixels with mostly constant surface colour.

score of zero, red being an anomaly score of 10 mm. As can be seen, most of the clear deposits present are identified by the anomaly detection method. In subfigure 5.5.3(a), all of the clear deposits on the left have been identified, and also some of the less obvious deposits on the right can be made out, for example on the low-right at half-depth. Additionally, some of the surface roughness is indicated at the top. Note also the large occlusion areas, which do not play a role in our detection algorithm, but further strengthen the identified anomalies. The pipe in subfigure 5.5.3(b) shows some deposits hanging from the top of the pipe, as well as considerable unevenness in the surface texture throughout the pipe that is also captured by the framework as local anomalies. In this pipe, however, we also see evidence of some false positives in the detection, in the lower far corners of the pipe. It appears that here, the point cloud contains too little data in the lower regions (closer to the camera, the field of view does not contain the bottom of the pipe), such that the model parameters are possibly inaccurate. In this particular area, there indeed appear to be some deposits around the flood line, but not to the extent indicated by the model. This phenomenon, that also plays a small role in subfigure 5.5.3(c), appears to be an artefact of the (in hindsight somewhat unfortunate) choice of lens that doesn't allow a full view of the pipe. The problem is easily remedied by removing the near and far end of the pipe and focussing on the middle band of the point cloud, which incidentally also concerns the pixels with the best focus. Remember that in *in-situ* inspections, the camera will be slowly inched forward through the pipe, allowing for a complete sweep of the pipe. Our framework can thus focus on the band of data where results are the most reliable.

Of the 48 image sets, a total of six images could not be fitted properly with our RANSAC method (not reaching the required fraction of inliers). Five of these concern the cases mentioned in the previous section as suffering from

a poor stereo matching. The resulting point cloud has a non-negligible number of points that erroneously lie inside the pipe, reducing the number of inliers. The sixth pipe that could not be fitted properly had an entire section broken off, presumably during extraction from the soil. The remaining 42 cases (87.5%) were properly fitted, producing the marked point clouds demonstrated in figure 5.5.3, as well as a single global anomaly score per image, as defined in equation (5.14). Prior to validating the anomaly detection, the photographs of the pipes were also graded subjectively in terms of defect severity, to have an independent ground truth to compare the detected anomalies levels with. For the 42 cases where our framework produced a global anomaly score, a Pearson's correlation of $r = 0.65$ with the ground truth was obtained, which is, according to the customary interpretation, a moderate, positive correlation.

5.5.3 DISCUSSION


Although our framework demonstrates good results on the two defect types of exposed granulate and deposits, not all pipes are correctly assessed. The main weakness of our method appears to revolve around the images with smooth pipe surfaces. The problem with these pipes occurs in the first part of our framework, the stereo matching, which indeed is known to be problematic in images with limited texture. The upside of this limitation is that pipes with smooth surfaces (often fairly new pipes) typically do not contain any defects. The main problem here is hence one of false alarms: the method sometimes erroneously identifies defects in smooth pipes because points are incorrectly placed in the 3D space. In future work, we may investigate whether stereo matching could also produce a confidence score, indicating the quality of the stereo matching in each region of the image. If successful, the method would only

identify a defect if both the confidence is high (in other words, not a smooth surface) and the anomaly score is high.


Whenever the stereo matching produces at least a reasonable result, the surface fitting and anomaly detection correctly identifies the various defects present. It should be noted that the framework even identified defects that were overlooked in the initial subjective quality grading of single images (not stereo pairs). In that sense, other than merely automating parts of the inspection process, our framework also has the potential to outperform human inspectors in certain respects.

Our experiments were performed in the lab, which may have had some minor effects on the outcome, although both in a positive and a negative sense. On the positive side, our experiments were perhaps made more challenging than necessary due to some initial choices that ended up being suboptimal. For example, the chosen lenses were not of sufficiently wide angle, such that the pipe could not be captured within the region of focus entirely, especially for egg-shaped pipes. The effect of this on the results is that for some regions of the pipe, such as the corners of the point clouds corresponding to the edges of the image, not enough evidence of the regular geometry is available in order to reliably decide on deviations from that geometry. This effect, which may cause both false positives and false negatives locally, can be observed in subfigure 5.5.3(c). This problem can be easily corrected by using a lens with a wider field of view. Another unfortunate choice concerns the low lighting conditions, which could have been corrected with a longer exposure of the images. Low lighting has not played a significant role (as the mostly good results demonstrate), but may have contributed to the lack of matching in areas with smooth surfaces.


Another property of the lab set-up was that we search for anomalies in the entire pipe, which is not necessary in the field, and produces some challenges with limited focal depth and missing parts of the pipes (due to extraction). In




(a)



(b)



(c)



(d)

Figure 5.7: Examples of four anomaly scores. Left shows the reference image, right shows the identified anomalies. Note that the anomalies on the edges and outside of the pipes are ignored.

a more practical setting, in an actual *in-situ* pipe inspection, a single image pair would only be inspected for a somewhat narrow band, corresponding to the region of focus. After that, the pipe inspection vehicle would move forward by a small distance and the process would be repeated. Note that although our framework makes no such assumption, being in pipes of mostly the same geometry would allow one to assume a certain steady geometry, and deviations from this could be more easily recognised.

On the negative side, our lab set-up is less realistic as we only inspect individual pipe segments, not longer pipe systems. As a result, we have no images of pipe joints, which in most cases would produce a (false positive) deviation from our fitted surface. Although we expect the pipe joints to be easily matched by the stereo matching, and proper handling of these ‘acceptable anomalies’ would not be difficult, our current method does not include such facilities, nor have we been able to test this. As future work, it would be interesting to develop a method that uses 3D point clouds of joints in order to recognise joint-related defects such as misaligned joints or signs of leakage. A final difference between our setting and actual sewer systems is that the inspected pipes, after having been washed out, will be wet, whereas our lab pipes were recorded in a dry state. We do not expect this difference to have a significant effect on the results, but would need to set up new experiments to assess this. The experiments were performed prior and during the COVID-19 lock-down, and some of the pipes were subsequently part of destructive full-scale testing. As such, we cannot easily redo the experiments.

In future experiments it may be interesting to look at discontinuities in the points on the pipe surface as well, as these may be caused by occlusions. An occlusion of a portion of the pipe surface may be just as informative as the geometry of the visible parts of the pipe surface, but could be indicative of foreign objects present in the pipe, such

as roots. However, because poor stereo matching may also result in such discontinuities, it may be advisable to require that such discontinuities exist only in the point cloud, and not in the disparity map.

5.6 CONCLUSION

5.6.1 SUMMARY

In this chapter we have proposed RADIUS, an anomaly detection framework for sewer pipes based on computer stereovision. The framework consolidates several successful techniques into a sequential process, to allow for anomaly detection in an automated fashion from stereo photographs without intermediate user input. We performed experiments to demonstrate the efficacy of the framework and conclude that it is successful in detecting defects present in physical pipes as anomalies in the three-dimensional geometry, and moves the state of the art closer towards fully automated sewer asset management.

5.6.2 LIMITATIONS

The major limitation of this work is the varying quality of the data obtained in the lab, as a result of our limited experience with the hardware. Some of the images were made in poor lighting conditions and without proper camera calibration. Repetition of the experiments was not possible due to time and budget constraints, and because a part of the pipes had been subjected to destructive full-scale testing in other experiments. A secondary limitation related to this is the total number of experiments performed.

A more intrinsic limitation of the approach of stereovision is that smooth, undamaged concrete may not contain

enough texture or markers to accurately match the reference image to the secondary image. While this is potentially an issue for a large portion of sewer pipes, we argue that such regions are unlikely to contain any anomalies. That said, since the absence of a match may also be an indicator of occlusion, we advise authors of future research to distinguish causes of a lack of a match: in the case of a too smooth pipe, the cause is likely a match that does not meet the uniqueness ratio required, as opposed to a patch in the reference image that does not appear in the secondary image due to occlusion.

In spite of these limitations, we feel the efficacy of the framework has been more than adequately demonstrated.

5.6.3 RECOMMENDATIONS

While anomaly detection may be a goal in itself, we hold the view that it is a stepping stone towards fully automated sewer condition assessment. To this end, we recommend future research to be performed into a follow-up step for the proposed framework: automated classification of the anomalies into defect classes. We feel that (the deviation from) the surface found through robust regression has a lot of potential for classification, as it will theoretically contain very little noise, as well as have a notion of “expected” behaviour.

It must be noted that while we have shown stereovision to be a viable tool for sewer pipe defect classification, the added value in practical settings has yet to be demonstrated. We have designed this framework to be compatible with current inspection practices: (monovision) CCTV inspection can still be performed while collecting data from two camera sources for stereovision experiments. This data may be used in parallel in order to both inch the industry towards automation of inspections, as well as to improve manual inspection techniques with an additional mode of data. State-

of-the-art sewer defect detection solely based on CCTV data may suffer from a relatively large false positive rate³⁴, but the additional depth information provided by an additional camera could lower this significantly. While the ambition of automated inspection is currently en vogue (again), the added value of multi-sensor inspection for more reliable, precise, and complete detection of a range of observable sewer defects, is an important added value worth researching further.

³⁴ MEIJER, D. W., SCHOLTEN, L., CLEMENS, F. H., AND KNOBBE, A. J. 2019. A defect classification methodology for sewer image sets with convolutional neural networks. *Automation in Construction* 104, 281–298