Machine learning and computer vision for urban drainage inspections

Meijer, D.W.J.

Introduction

1

This thesis aims to expand the available knowledge on how machine learning and computer vision techniques can be used to improve efficiency and quality of urban drainage inspections. This chapter will outline the motivation for the research and the contents of the rest of the thesis.

1.1 Motivation

1.1.1 Sewer Asset Management

Properly operating urban drainage systems are essential to ensure public health, safety, and productivity in cities, but not enough is known about the failure mechanisms that lead to decreased performance or loss of functionality [1]. To understand the condition of the system and to assess which assets need repair or rehabilitation, inspections are performed.

The largest share of the operation and maintenance cost across the technical assets in the system is usually spent on the sewer pipes. For their inspection, CCTV inspection is commonly performed: a 'pipe inspection vehicle' (*PIG*) is lowered into a manhole, where it records photo or video footage which is reviewed by trained operators. The operators identify defects and possible indications of defects in the footage, and assign this a severity rating between 1 (no intervention necessary) and 5 (immediate intervention necessary).

One of the major shortcomings of this method is that these severity ratings and the defect identification prior to it are highly subjective, and have been shown to differ not only between operators, but also for the same operator at different time points [2] [3].

[1] Stanić, N., Langeveld, J. G., and Clemens, F. H. 2014. Hazard and operability (hazop) analysis for identification of information requirements for sewer asset management. *Structure and Infrastructure Engineering 10,* 11, 1345–1356

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

[3] Wirahadikusumah, R., Abraham, D., and Iseley, T. 2001. Challenging issues in modeling deterioration of combined sewers. *Journal of infrastructure systems 7,* 2, 77–84

A second shortcoming is that these urgency ratings assigned by operators do not necessarily reflect the actual urgency of intervention in a broader sense. It is only a measure of how advanced a specific defect appears to be, and does not take into account any other factors, such as location of the pipe or co-occurrence with other defects. As such, it is not a viable measure of the risk or costs incurred by neglecting to intervene.

1.1.2 Machine Learning and Computer Vision

At the start of the SewerSense project [4], recent advances in machine learning and computer vision had for the most part not yet been introduced to the field of urban drainage. The SewerSense project aimed to incorporate such recent advances to the sewer inspection task.

[4] TISCA programme funded by NWO-TTW. 2016-2020. Sewersense – multi-sensor condition assessment for sewer asset management

Machine learning techniques such as neural networks and other classification algorithms can potentially automate parts of the inspection by processing the photo and video footage faster and more precisely than a trained operator could. Automating parts of the inspection process can on the one hand objectify the inspection results, and on the other hand facilitate decision making in sewer asset management by providing more accurate information than an arbitrary urgency scale.

In the short term, these techniques may also be integrated into current inspection practices to increase the inspection efficiency and quality, while the industry slowly moves toward fully automated inspections. The visual inspections that are the current practice can be used to train computer vision algorithms that can extract knowledge from images, which can bring short term benefit even if the eventual fully automated solution does not rely on visual inspection.

Full automation may become possible in the future, but will require more and higher quality data than is available at the time of writing this thesis. Some defect types are exceedingly rare, difficult to spot, and experts may not agree on them. Expecting an algorithm to accurately detect such defects is currently not realistic. In the long term, however, trained machine learning and computer vision algorithms may reshape the sewer inspection practices to rely less on human inspections.

1.1.3 Scope

The scope of this thesis then, is to perform preliminary research into the possibilities of machine learning for automation in the asset management industry. We attempt to bridge the gap between current sewer inspection practices and state-of-the-art machine learning and computer vision techniques that promise to automate (parts of) the process. We explore how well machine learning and computer vision algorithms can perform on existing data from previous inspections, as well as how an additional mode of data collection can improve this performance while remaining compatible with current inspection practices.

This thesis will not cover the collection of visual inspection data that is used as input for our models, nor the decision making process of when and whether to repair or replace. These aspects are both covered extensively in their respective domains.

1.2 Research Questions

We pose several research questions to be answered in this thesis. These question roughly correspond with chapters of the thesis, and the answers will be summarised in the conclusion.

Q1

WHAT KNOWLEDGE CAN BE OBTAINED FROM AVAILABLE INSPECTION DATA WITHOUT THE UTILIZATION OF EXPERT CLASSIFICATION, WHICH MIGHT BE INCONSISTENT OR UNAVAILABLE?

We know expert classification to be limited in reliability [5], and much data collected from inspections may not have the expert classifications available in machine-readable format. Unsupervised learning techniques can obtain knowledge from this data simply by clustering it in similarity to itself, and may circumvent these two issues altogether.

[5] DIRKSEN, J., CLEMENS, F., KORVING, H., CHERQUI, F., LE GAUFFRE, P., ERTL, T., PLIHAL, H., MÜLLER, K., AND SNATERSE, C. 2013. The consistency of visual sewer inspection data. *Structure and Infrastructure Engineering 9*, 3, 214–228

Q2

HOW CAN THE DATA COLLECTED WITH CURRENT INSPECTION PRACTICES BE ANALYSED WITH MACHINE LEARNING TECHNIQUES IN ORDER TO IMPROVE PROCESSING EFFICIENCY AND ACCURACY?

While fully automated sewer inspections might not be a pipe dream, a more timely benefit may be found in incremental progress. Instead of rebuilding the industry from the ground up for automation, we investigate whether analysis of existing data with modern computation can improve quality today.

Q3

HOW DO WE ASSESS THE QUALITY AND OPERATIONAL IMPACT OF (PARTIAL) AUTOMATION OF THE CURRENT INSPECTION PRACTICES?

Training a model to make predictions is of limited use if we have no reference to compare the modelled predictions to reality. For such a model to be used in practice, we must have an assessment framework with which to assess its performance. Different types of errors will have differing consequences, and the impact on the industry might not be measured by metrics that are common in machine learning.

To what extent are the current inspection practices automatable? Q4

Q1, Q2, and Q3 focus on building and verifying models from existing data, we must also look beyond the current practices and investigate the limits of such approaches. Restricting models to the existing data may hinder progress or reinforce existing biases from current practices. It is important then, to decide whether building on the current practices is worthwhile, compared to a bottom-up design approach focused on automation entirely.

Does introducing depth information through computer stereovision improve the data quality and analysis capabilities? Q5

Extending Q4, we investigate a specific additional mode of data collection, that of stereovision, the use of two cameras to obtain not only a two-dimensional image, but a depth component as well. Such advanced modes of data collection are not exactly novel, but have not seen much use in practice because of the additional training required for human operators to interpret the results. From our scope of preparing the industry for automation, it is then again an interesting question to see what the added value of such data modes is.

How can we employ machine learning and computer vision to improve the efficiency and quality of urban drainage inspections? Q6

Using the answers to the previous five questions, we conclude having outlined the possibilities of enriching sewer asset management with machine learning and computer vision techniques, and highlight the areas that still require more research.

1.3 CONTRIBUTION

Many of the ideas and text in this thesis have appeared earlier in four publications, written over the course of five years. These papers are an extended abstract in a regional conference, a continuation of that extended abstract in the proceedings of an international conference, and two articles published in a high impact journal.

This section presents an overview of each of the main chapters in this thesis, and lists the main contributions made for that chapter and in which publication the contributions first appeared.

CHAPTER 2 goes over preliminary knowledge required for a complete understanding of this thesis.

CHAPTER 3 approaches the defect detection problem from an unsupervised learning perspective, based on how commonly patterns appear in a set of images. The approach leverages principal component analysis or a convolutional autoencoder for their ability to generalise only those patterns that appear frequently in the training set, interpreting poor generalization as a signal of anomalous information. The main contribution is the comparison of the original image to an image partially reconstructed by the autoencoder or a limited number of principal components. Parts of this chapter have previously appeared in [6] and [7]. This chapter explores research questions Q1 and Q4.

[6] MEIJER, D. W. AND KNOBBE, A. J. 2017. Unsupervised region of interest detection in sewer pipe images: Outlier detection and dimensionality reduction methods (extended abstract). In *Benelux Conference on Machine Learning (BeneLearn)*

[7] 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

CHAPTER 4 approaches the defect detection problem from a supervised learning perspective, using a convolutional neural network. The article this chapter is based on became highly influential, having been cited over two dozen times at the time of writing this thesis. This might be attributed to timeliness or novelty, but we feel that the main contribution is the methodological groundwork that explores the results as relevant to the target domain: the realistic conditions in which the experiment was performed, to ensure a realistic assessment of real-world value of such a classifier. Parts of this chapter have previously appeared in [8]. This chapter explores research questions Q2, Q3, and Q4.

[8] 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

CHAPTER 5 introduces a new data modality by using two cameras to capture sewer pipes in stereovision. The added depth channel is combined with knowledge of the physical properties of the setup to reconstruct the three-dimensional pipe geometry virtually. Anomaly detection is performed through robust regression of a model that is informed by the expected geometry of a sewer pipe, under the assumptions that certain types of surface damage will deviate from this model. The main contributions are the adaption of stereovision techniques to this unique use case of an object that is positioned perpendicular to the image plane, and the sewer pipe model. Parts of this chapter have previously appeared in [9]. This chapter explores research question Q5.

[9] MEIJER, D. W., LUIMES, R. A., KNOBBE, A. J., AND BÄCK, T. H. W. 2021. RADIUS: Robust anomaly detection in urban drainage with stereovision. *Automation in Construction 139*, 104285

CHAPTER 6 concludes with a summary of the main content of the thesis and answers to the six research questions.