

Grouping of Maintenance Actions with Deep Reinforcement Learning and Graph Convolutional Networks

This work addresses the problem of optimal maintenance planning for a sewer network using machine learning techniques. We propose a Deep Reinforcement Learning framework based on Graph Convolutional Networks to leverage the graph structure of assets. Instead of planning for individual pipes, we group maintenance on spatially close assets. The framework is evaluated on a sewer pipe network and the results are compared to several baselines. Our approach shows potential for developing efficient and practical maintenance plans in terms of cost and reliability.

Introduction

This project was performed in collaboration with Rolsch Assetmanagement, a software development and consultancy company in the field of underground water and sewer pipes.

We address the problem of optimal maintenance planning using historical data and propose a framework based on two machine learning techniques: Deep Reinforcement Learning (DRL) and Graph Convolutional Networks (GCN). Reinforcement learning is the problem of learning behaviour through trial-and-error interactions with a dynamic environment. The goal of DRL is to learn an optimal policy for decision problems by maximizing a cumulative reward, making use of a neural network. We use a graph-based neural network architecture (GNN) to leverage the graph structure of pipe assets, inherently present due to the physical connections between them. GCN is a type of GNN that defines convolution for graph-structured data.

As a demonstrator, we employ a sewer pipe network, which is an essential part of urban infrastructure. Failure of pipe assets can cause service disruptions, public health threats and damage to surroundings. Because the sewer infrastructure is underground, inspections and rehabilitation are expensive and labor-intensive, while the budget is often constrained. Therefore, a maintenance strategy balancing reliability and costs is needed to achieve an adequate level of service.

We investigate the capabilities of DRL to find the best rehabilitation moment for groups of

pipe assets in terms of cost and reliability. We try to reduce the cost by considering the physical state of the neighboring pipes when choosing rehabilitation actions. Grouping of maintenance can save the additional setup, labor and unavailability cost of the network. The objective of the study is to show the potential of the DRL framework for solving maintenance planning problems on a network of assets.

Problem

There are a number of assets that may deteriorate over time. Each asset has a certain status related to the deterioration that depends on age and other properties. Based on its status, each asset may need one out of multiple possible maintenance actions. There is a cost associated with the maintenance actions and there is a cost imposed in case an asset is near failure because of a lack of maintenance. The assets form a network, i.e. they are structurally connected. We are interested in the simultaneous rehabilitation for groups of geographically close assets instead of interventions on individual assets at different moments. The overall objective is to plan maintenance such that the assets do not deteriorate to near failure while the overall cost is minimized.

Case study: a sewer pipe network

As a driver case for our work, we consider a case study of a real sewer pipe network of the city of Breda with a total length of 2170 km, comprising pipe properties and inspection records. We choose a subset of 942 pipes (38 km) to evaluate the DRL framework. The DRL agent implements a Deep Q-Network, which is

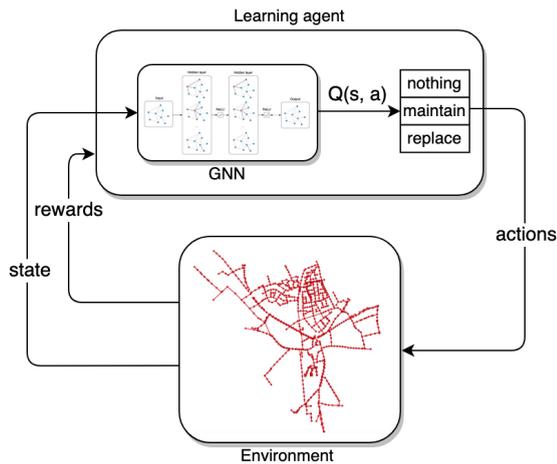


Figure 1: Interaction between DRL agent and environment, and the position of the GCN within the agent.

a reinforcement learning algorithm that incorporates a neural network to learn a policy. The neural network used is a GCN. The interaction between the DRL agent and the environment is depicted in Figure 1. In order to apply a GCN, the data is represented by a graph. Each node corresponds to a sewer pipe and there is an edge between nodes if their pipes are within a given distance in the real world.

Simulation environment

Pipe deterioration is modeled using the exponential distribution, which is commonly used for modeling the lifespan of deteriorating assets. For each pipe a failure rate is extracted from a dataset of manual pipe inspections. Each node has a feature vector including *length*, *material*, *age* and failure probability pf . Each time step one of three actions is chosen for each pipe: *do nothing*, *maintain* and *replace*. The *age* represents a pipe's physical state and is updated each time step to reflect the applied action. If *do nothing* is applied to a pipe, the *age* increases by 1 year, if *maintain* is applied, it decreases by 10 years and if *replace* is applied, it is reset to 1. The pf is updated based on the new *age*. Performing *maintain* and *replace* actions incur a cost and if intervention is applied on two spatially close pipes, a small cost reduction is applied to both to encourage grouping.

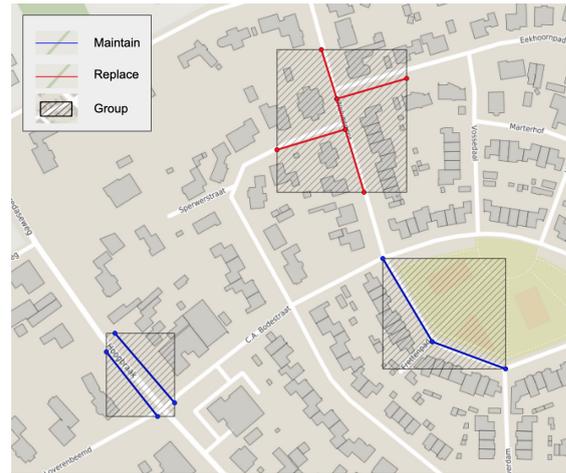


Figure 2: Example of a grouping on pipes that are close to each other, produced by the GCN approach.

Results

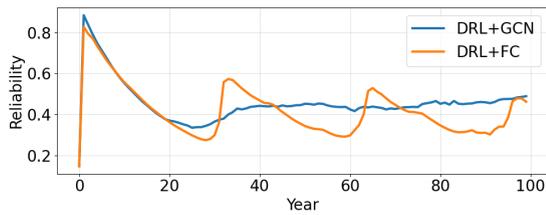
After training, we generate a maintenance plan with a 100-year time horizon for the GCN approach and four baselines. The baselines are: a simpler variant with only a fully connected layer instead of GCN (DRL+FC), greedy preventive (*maintain* when $pf > 0.5$), time-based preventive (*maintain* every 10 years) and corrective (*replace* when $pf > 0.95$). The goal is to achieve intervention grouping to optimize maintenance plans. The higher the number of pipes per group, the more rehabilitation is concentrated in a smaller amount of different geographical locations. The GCN creates a plan in which 56% of the groups have more than one pipe, resulting in an average of 2.74 pipes per group. See Figure 2 for an example of such grouping. A summary of the results and comparison with baselines is shown in Table 1. Figures 3 and 4 show respectively the reliability and cost of each approach. We see that although the GCN strategy is slightly more expensive than DRL+FC, it provides better reliability, a higher degree of grouping, and lower number of interventions.

Conclusion

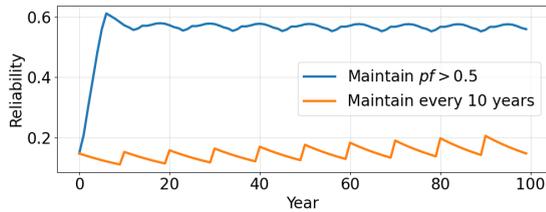
This work presents a DRL framework that combines DRL and GCN for the rehabilitation planning of sewer pipes. The DRL agent learns an improved policy in terms of lower cost and

	DRL+GCN	DRL+FC	Replace	Greedy	Maintain-10
Mean # pipes per group	2.74	2.02	-	-	-
% groups with >1 pipe	56%	30%	-	-	-
Mean reliability	0.46	0.42	0.32	0.56	0.15
Interventions per year	45.26	55.94	20.77	378.45	94.2
Interventions per pipe	4.80	5.94	2.44	40.18	10
Total cost based on reward function	8362	8077	21,769	15,526	50,908

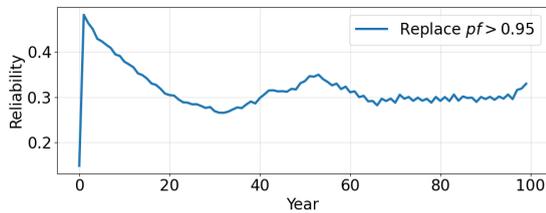
Table 1: Comparison of DRL+GCN with: DRL+FC (fully connected layer), one corrective and two preventive baselines.



(a) DRL approach.



(b) Preventive approach.

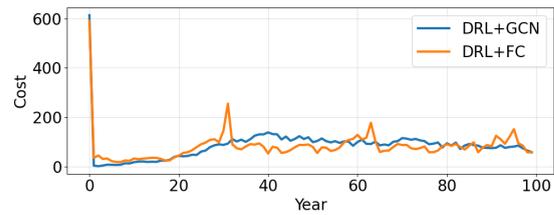


(c) Corrective approach.

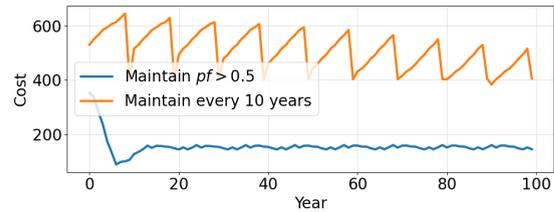
Figure 3: Mean reliability per year for each strategy.

higher reliability and uses GCN to leverage the relational information encoded in the graph structure of the sewer network. Our framework is successfully evaluated on a real dataset to show its potential for applications in infrastructure maintenance planning. The proposed approach is not intended to solve the specific case study described in this work but to serve as a feasibility study for applying the combination of DRL and GNNs for asset management problems.

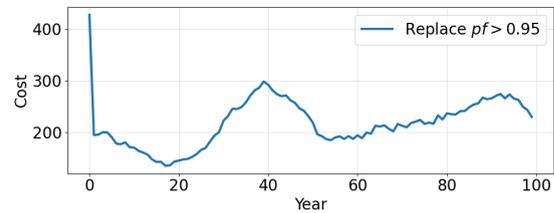
Although the GCN approach doesn't result in the lowest cost in the used configuration, it



(a) DRL approach.



(b) Preventive approach.



(c) Corrective approach.

Figure 4: Total cost per year for each strategy.

provides good overall reliability, a structurally higher degree of maintenance grouping and less interventions in total, compared to the baseline methods. It also shows a steady performance of the sewer network across the 100-year time horizon, in terms of annual costs, number of interventions and reliability. The results show that there is potential for a DRL-based approach incorporating GNN for maintenance optimization problems in the domain of asset management, and they provide motivation to further explore this direction.



Facts

Student David Kerkkamp

University Radboud University

Supervisors Zaharah A. Bukhsh
Nils Jansen

Company Rolsch Assetmanagement

PrimaVera - Powered by



**nationale
wetenschaps
agenda**