Predicting Water Pipe Failures with a Recurrent Neural Hawkes Process Model

Citation for published version (APA):

Verheugd, J. T., de O. da Costa, P. R., Refaei Afshar, R., Zhang, Y., & Boersma, S. (2020). Predicting Water Pipe Failures with a Recurrent Neural Hawkes Process Model. In *2020 IEEE International Conference on Systems, Man, and Cybernetics, SMC 2020* (pp. 2628-2633). [9282941] Institute of Electrical and Electronics Engineers. <https://doi.org/10.1109/SMC42975.2020.9282941>

DOI:

[10.1109/SMC42975.2020.9282941](https://doi.org/10.1109/SMC42975.2020.9282941)

Document status and date:

Published: 14/12/2020

Document Version:

Accepted manuscript including changes made at the peer-review stage

Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

www.tue.nl/taverne

Take down policy

If you believe that this document breaches copyright please contact us at:

openaccess@tue.nl

providing details and we will investigate your claim.

Predicting Water Pipe Failures with a Recurrent Neural Hawkes Process Model

1st Jeroen Verheugd
School of Industrial Engineering
Eindhoven University of Technology
Eindhoven, The Netherlands
jeverheugd@deloitte.nl

2nd Paulo R de Oliveira da Costa
School of Industrial Engineering
Eindhoven University of Technology
Eindhoven, The Netherlands
p.r.d.oliveira.da.costa@tue.nl

3rd Reza Refaei Afshar
School of Industrial Engineering
Eindhoven University of Technology
Eindhoven, The Netherlands
r.refaei.afshar@tue.nl

4th Yingqian Zhang
School of Industrial Engineering
Eindhoven University of Technology
Eindhoven, The Netherlands
yqzhang@tue.nl

5th Sjoerd Boersma
Vitens N.V.
Zwolle, The Netherlands
sjoerd.boersma@vitens.nl

Abstract—Water distribution networks have shown an increased rate of failure due to material deterioration. In this paper, we apply a Recurrent Neural Hawkes Process model to learn the failure intensity function of water pipes. The failure intensity function is learned based on two components: the base failure rate that is determined by the unique pipe profile attributes, and the effect of past failures. Compared to the existing solutions, our model is able to predict the time to next failure on an individual water pipe level. The learned failure intensity function is used to identify value points in the deterioration process of water pipes that represent their economical end-of-life. We use data from a Dutch water distribution network that consists of 49,600 km of pipelines to test the performance of the proposed model. We have made this dataset available online.

Index Terms—Predictive maintenance, Water pipe failure, Intensity function, Point Process, Recurrent Neural Networks

I. INTRODUCTION

The water distribution network in the Netherlands has slowly expanded since the 19th century. Currently, installed water pipe lines are approaching the end of their expected life time and have started failing at an increasing rate. Water companies monitor the quality of their water distribution network and prevent water pipe failures to guarantee drinking water standards. The currently operational water pipes use four main types of material: cast iron, asbestos cement (AC), polyvinylchloride (PVC) and polyethene (PE). Each material has its own characteristics, which results in different estimated life expectancy. This, in combination with various structural sizes and environmental effects makes the distribution system prone to water pipe breaks and leaks [1].

In the water distribution practice, most companies evaluate their existing infrastructure by rating the water pipes on an index from excellent to inferior quality condition [2]–[4]. This index allows for prioritization in a water pipe replacement program. However, time is not taken into consideration, which makes the index inappropriate to predict the timing of future water pipe failures. To overcome that, survival analysis is a

commonly used approach for predicting water pipe breaks within a certain time period. It is normally applied to homogeneous groups of water pipes based on pipe age, material type and soil type characteristics [5]. However, this approach does not allow to accurately determine the probability of failure on an individual water pipe level.

Existing studies [6], [7] have suggested that the failure intensity of an asset, i.e. its hazard rate, normally stays at a relatively stable level. This can be regarded as a base component depending on its profile covariants. Moreover, the occurrence of a failure can often lead to an instantaneous rise of an asset's vulnerability. The reason behind this is that when a water pipe failure occurs, it often becomes more fragile to failures due to the fundamental physical damage. This vulnerability gradually fades back to the baseline when the asset recovers e.g., after maintenance. Lastly, different types of failures have different triggering effects to each other, e.g., a pipe burst will cause more damage to a leak failure [8].

This implies that failure events can be correlated. The ability to discover correlations among events is crucial to accurately predict the future of a sequence given its past, i.e., which events are likely to happen next and when [9]–[11]. An event sequence, like failure events of water pipe lines, carry important clues about the underlying dynamics, and is different from time-series sequence whereby a series is indexed by fixed time intervals. Rather, this event sequence is often associated with continuous time stamps and additional information such as event type or profile features [8]. A suitable framework for modeling events is a *point process*, a mathematical model of phenomena represented as occurring events in a space.

In this work, the failure intensity function of water pipes is learned by combining recurrent neural networks with point process modeling [8]. Contrary to current solutions proposed to model water pipe failures, our model is able to predict the time to next failure on an individual water pipe level. Our method is believed to be expressive enough to learn

the underlying dynamics of water pipe failures without a predefined parametric function.

The main contributions of this work are: (1) We model the failure intensity function of water pipes via a base intensity rate and long-term effects of historical failure events, respectively determined by unique profile attributes and failure event sequences. This approach results in better performance of the time to next failure than a function learned only using the profile attributes or failure events; and (2) We incorporate novel effects to the profile attributes of water pipes by including the distance to the closest water production location, i.e. water hammer effect.

II. RELATED WORK

One line of work investigates the effect of physical factors on the deterioration of water pipes. In [1], a case study analysis identifies that attributes such as diameter, length, age, material, soil type, and failure history are important factors for predicting water pipe failures. In [12], the buried depth of water pipes is included as an important physical parameter. In addition, the effect of tree root intrusion on nearby water pipe lines has been captured in [13]. Another important indicator of water pipe failures is water pressure. However, this operational parameter is difficult to measure. In [5] the authors categorize each water pipe in a water pressure zone due to the elevation differences of the city. A common approach is determining the nominal water pressure with hydraulic models, see [2], [12], [14], [15]. In [16], results imply that the effect of water pressure on failures is moderate, however, high pressure peaks have a larger influence on water pipe deterioration. Another water pressure related effect is called water hammer, which can be best explained as a sudden water pressure surge caused when water in motion is forced to stop or changes direction. To the best of our knowledge, no existing work studied the effect of water hammer on water pipe failures. We investigate this factor in our prediction model.

As previous failures can have a temporal effect on the failure rate of an asset, one can use the Hawkes process [17], a type of point process, to model this underlying behavior. In [3], a profile specific self-exciting Hawkes process is proposed to model the past failure events in combination with a profile specific base intensity. However, the proposed model requires prior domain knowledge in the form of a parametric failure rate function and does not allow for a time-varying base intensity rate. These limitations are reduced when the intensity function is modelled with a Recurrent Neural Network (RNN) in [8], where the background and history effect of point process are learned using two RNNs on a maintenance case of ATMs. In our approach, we extend the idea in [3], [8] and employ an event sequence RNN to capture the long-term effects of historical water pipe failures, such as exciting and inhibiting effects. In addition, we embed a Neural Hawkes Process model that allows to model the background intensities and long-term effects of past failures per individual instances of water pipes, for a major Dutch water distribution network.

III. DATA PREPARATION

To model the intensity function of water pipe failures there are two relevant sources of data, i.e., distribution network assets and historical failure records. Several data transformation steps are taken to be able to model the failure intensity function with the Neural Hawkes Process model. Furthermore, the water profile attributes are enriched with operational, environmental attributes and external geographical charts.

A. Water distribution network data

The water distribution network of our industrial partner consists of more than 1.2 million unique water pipe IDs. Only a selection of water profile attributes within this dataset are relevant for our predictive maintenance problem.

1) *Data collection*: Our industrial partner has provided data of all water pipes that were in operation as of June 2019. However, water pipe renewal records are not present. The failure events that have occurred on these replaced water pipes are important data that give insights in the deterioration process of water pipes. Therefore, two older archived datasets, respectively 2011 and 2017, have been used to determine which unique water pipe IDs have been replaced between 2011-2019. This resulted in 1,212,205 operational assets in 2019; 6,784 expired assets between 2017 - 2019 and 66,885 expired assets between 2011-2017.

2) *Data Quality*: After validating with the domain experts, only water pipes with a diameter between 35 - 800, installation date between 1857 and 2019 and pipe lengths above 1m are considered. 13.2% of pipe IDs had unknown installation dates. Therefore, a public dataset including all dwellings in the Netherlands along with their construction year [18], is used to retrieve the installation date for these pipes. The assumption made here is that most water pipes are installed during the same construction period of the nearby buildings.

Finally, each unique water pipe ID is linked to a geographical location, denoted as geometry.

B. Failure records of water pipes

Historical failure records give insights in the life cycle of water pipes in the distribution system. In the period of 2005 till 2019 there are 35,347 registered historical failures, with a corresponding cause of failure, year of failure, type of failure, diameter of the water pipe and GPS coordinates of the location. From the total number of water pipe failures recorded, the majority of failures are related to the ageing process of the water pipes, fractures due to soil subsidence and corrosion of the material type, as shown in Figure 1.

Fig. 1: The total number of failure events per type of cause

Not all failure events are recorded with an accurate timestamp of occurrence. On average 2,000 failure events occur per year. However, the data becomes of quality in the year 2009 with only a few failure events between 2005-2009. This has led to a total number of 19,579 considered failures. The majority of failures have occurred on water pipes of the material types AC, PVC and CI. There is a lack of sufficient failure events of the other material types to truly understand their failure pattern and are therefore excluded in further analysis.

C. Enrichment of the datasets

1) *Creation of asset units.*: Some water pipe IDs share the same profile attributes, are physically connected and have been installed on the same date. In theory, these water pipes have the same age, experience the same local environmental effects and therefore have a similar deterioration process. These groups of water pipe with shared profile attributes are from now on referred to as asset units. The creation of asset units is fundamentally different than homogeneous groups of water pipes based on the requirement of a physical connection between them. Based on the domain knowledge available, water pipes are grouped based on the following characteristics: (i) the asset unit shares the same material type; (ii) the asset unit has the same diameter; (iii) the asset unit has been installed in the same year and (iv) all water pipes in the asset unit must be connected within a buffer of 16 meters. To compensate for the partial replacement of pipelines within an asset unit, a buffer of 16 meters is used based on the standardized tube length.

Fig. 2: Visualization of created asset units

2) *Allocating failure events.*: Although most failure events seem to be located on top of the distribution network, for some of them it is unclear to which asset unit they are allocated (Figure 2). In areas with a high density of water pipes, a small inaccuracy in the GPS coordinates can result in a false allocation. To alleviate these errors, we develop a geospatial algorithm based on GeoPandas¹. Based on the characteristics of the failure event, i.e. material type and diameter, the algorithm allocates each failure to the nearest asset unit of the

same characteristics. As a measure of accuracy, the amount of distance moved is calculated. All failures moved within 100m are considered reliable, resulting in a total of 16,182 failures in scope for further analysis.

3) *Feature engineering.*: The effect of water hammer on water pipes, the number of appendages per 100 meter, the ground soil category and land vegetation above an asset unit is added to the profile. The effect of water hammer, measured by the distance to the closest water production location or water accelerator, is a new feature that has not been used for modelling water pipe failures in the literature. These production locations and accelerators can cause a sudden surge of water pressure, which could cause a water pipe to break. The number of appendages per water pipe is related to a higher risk of leaks. Each type of soil has a different effect on a specific material type, making it an important feature to include in the asset unit profile. The amount of vegetation near the location of a specific asset unit represents the risk of tree roots that can damage the structure of the water pipes. The land cover above an asset unit can be categorized into four classes, namely high green, low green, rural and non green. The amount of vegetation near the location of a specific asset unit represents the risk of tree roots that can damage the structure of the water pipes. A priority heuristic is used to maintain the highest vegetation class experienced per asset unit. A Spearman correlation analysis has indicated a strong non-linear correlation between the prediction label and our novel water hammer feature.

IV. RECURRENT NEURAL HAWKES PROCESS MODEL

The design of the Recurrent Neural Hawkes Process model consists of two components, a static vector that incorporates the profile features of the asset units, and a failure event sequence to learn the long-term effects of past failure events.

1) *Base Failure Intensity Rate.*: In the Neural Hawkes Process model the base intensity rate is learned from the profile attributes of the asset units. The profile attributes of the asset units consist of structural, operational and environmental characteristics that determine their unique deterioration process. The profile attributes material, diameter, function, length, appendages per 100m, distance to closest water pump, ground soil and vegetation of the asset unit are all static and do not vary over time. Categorical attributes in the static profile vector are one-hot encoded. The categorical profile attributes are transformed to 3 material, 3 function, 6 soil categories and 4 vegetation features. In addition, the numerical features are normalized using min-max normalization.

2) *Long-term effects of past failure events.*: The long-term effects of past failure events per asset unit are learned from a failure event sequence per asset unit. Two attributes of the historical failure records dataset are used for creating a failure event sequence, namely the timestamp of failure and the type of failure. These two attributes are used as the input for the event sequence RNN (an LSTM in our implementation).

The distribution of failure events is constructed in tuples with a type of failure and time of occurrence (k_1, t_1) , (k_2, t_2) ,

¹GeoPandas is a geospatial analysis package for the programming language Python

Fig. 3: Architecture of the Neural Hawkes Process model

..., where each $k_i \in \{1, 2, \dots, K\}$ is an event type and $0 \leq t_1 \leq t_2 \leq \dots$ are times of occurrence, denoted as number of years since installation. The idea behind such a construction is that the effect of past events on the base failure intensity rate differs per failure type. Past events may now either excite or inhibit future events. They do so by sequentially updating the states of the event LSTM.

All event sequences of the asset units start with a special beginning-of-stream (BOS) event (k_0, t_0) , where k_0 is a special event type and t_0 is set to 0. This special event type expands the LSTM's input by one and represents a normalized installation date of the asset unit, i.e. the birth event. The initial configuration determines the hidden state h_t and the intensity function $\lambda_k(t)$ over $t \in [0, t_1]$. Finally, the tuples in the failure event sequences must be transformed to enable the Neural Hawkes Process model to read these event sequences. Similarly to the static features, the event type is one-hot encoded to binary column values, and the times of occurrences are normalized to 0 – 1 via min-max normalization.

3) *Architecture*: The architecture of the Neural Hawkes Process model is visualized in Figure 3. The architecture consists of two input layers, an input layer for the static profile vector and a three-dimensional input layer for the event sequence, i.e., samples $\times$ event steps $\times$ features. The input of the event sequence layer is connected to the event sequence via LSTM cells. The static profile features and the output of the event sequence LSTM are merged into one vector and processed through two Dense layers. Finally, a timestamp prediction is made with a Dense output layer.

During training, the following parameters are used [8]: The state size of the event sequence LSTM is set to 16; the event set contains 6 time steps, namely a BOS event and the last 5 failure events occurred. After merging the static profile vector and the event sequence LSTM, the representations are processed through two dense layers of size 64 with softplus activation. A Dense layer of size 1 with a softplus activation layer is used as output. The Adam optimizer is used for learning the model weights during 60 epochs, learned with the Mean Squared Error (MSE) as loss function.

4) *Performance measurement*: The performance of the Neural Hawkes Process model (NHP) is determined by the

accuracy of predicting the timestamp of the next failure event. Three versions of the model are evaluated to understand the predictive performance of the individual components of the model: (1) NHP-P predicts the timestamp of the next failure event solely based on the profile features; (2) NHP-E predicts the timestamp of the next failure event solely based on the event sequences; and (3) NHP-C is the actual Neural Hawkes Process model architecture as shown in Figure 3. The performance of NHP-P and NHP-E are compared to that of NHP-C. During training of all model versions, the profile vectors and their corresponding failure event sequence will be split into a training subset and a test subset, respectively 75% and 25% of the total dataset. A 10-fold cross validation will be performed. The average prediction error of the 10 subsets is taken as the estimated prediction error. In addition to the loss function, i.e. Mean Squared Error, the Mean Absolute Error is calculated as an extra measurement.

V. EXPERIMENTS

The model variations of the Neural Hawkes Process model are learned in two different situations. (i) First, the model weights are learned based on a training dataset that included all asset units in scope. (ii) In the second situation, the training dataset is reduced so that it only includes asset units that have experienced at least two failures. It is believed that when the event sequence of the asset unit has at least one failure event besides the BOS event, the performance of the model improves. At last, the long-term effects of the failure event types are tested and evaluated. Performing these analysis has been performed on a Processor Intel(R)Core(TM) i7 CPU @ 2.80GHz with 16GB of RAM. The dataset and source code can be found online².

A. Target Label

The Neural Hawkes Process model uses the profile vector and failure event sequence to predict the occurrence, i.e. timestamp, of the next failure on an asset unit level. The units of this target label is denoted in the number of years since the year of installation. The predictions of NHP is used to compute the time difference between the predicted timestamp and the last failure occurred on that asset unit. Most asset units have not experienced a failure within our failure observation period. After excluding the assets units that do not have a target label, we have in total 10,203 units for experiments.

B. Performance

1) *Situation (i)*: In situation (i), the performance of the Neural Hawkes Process model is tested when the model is trained on all instances of asset units. Most event sequences of asset units only consist of one actual failure event, which will be used as the target label during training. For this reason, it is assumed that the predictive performance of the event sequence component in the **NHP-C** is limited in this situation. If the event sequence of the asset unit only consists of the BOS event, the Neural Hawkes Process model learns the failure

²<https://github.com/yingqianzhang/water-pipes-failure-prediction>.

Fig. 4: Distribution of probability of failure in situation 1

intensity function of an asset unit that is unconditioned on the historic failures of that asset unit.

In Table V-B1, it can be observed that the NHP-C, clearly outperforms its individual components NHP-P and NHP-E. The MSE of the NHP-C is 13.99 years with a Mean Absolute Error (MAE) of 8.79 years from the actual timestamp of the true label. In case of the NHP-P, the timestamp of the next failure is only predicted based on the base failure intensity rate determined by the profile features of the asset unit. The MSE of the NHP-P is 15.2 years on the training set and has a corresponding MAE of 10.5 years. The performance of the NHP-E represents the predictive ability of the long-term effects of past failures on the failure intensity function. The accuracy of the NHP-E is the lowest compared to the other two versions, a MSE of 17.6 years and MAE of 11.3 years. The low performance of the NHP-E can be explained due to the fact that most failure event sequences for training only consist of a BOS event.

To summarize the performance of NHP in situation (i), firstly, the combination of the base failure intensity rate and the long-term effects of past failures to learn the failure intensity function of water pipes have led to the best results. Secondly, the performance improvement of the NHP-C compared to NHP-P indicates that the Neural Hawkes Process model is able to distinguish whether it is predicting a first failure event or a consecutive failure conditioned on its history of failures.

The failure intensity function learned with the NHP-C is used to draw the probability of failure per material type, visualized in Figure 4. In these figures, the colors represent the confidence intervals σ , 2σ , respectively from dark to light blue. The red line represents the actual distribution as observed and the black line represent the distribution predicted with NHP. The NHP model is able to distinguish the distributions that differs per material type.

2) *Situation (ii)*: In situation (ii), the performance of the Neural Hawkes Process model is tested when trained on all instances of asset units that have experienced at least

Fig. 5: Predicted time to next failure of the Neural Hawkes Process model

two failure events. Contrary to the first situation, the failure intensity function of all asset units is now conditioned on past failures of the asset unit, which is incorporated in the event sequence. In total 10,203 asset units have experienced at least one failure, of which only 2,773 asset units are in scope in situation (ii). The performance of the Neural Hawkes Process model in situation (ii) is presented in Table I.

TABLE I: Performance of the Neural Hawkes Process model in situation (ii)

Data	Metric	NHP-P	NHP-E	NHP-C
Train	MSE	11.78	2.23	2.15
	MAE	8.12	1.79	1.59
Test	MSE	12.55	2.34	2.17
	MAE	8.60	1.78	1.64

The performance of the NHP-C has improved significantly to a MSE of 2.17 years and MAE of 1.64 years. Compared to situation (i), the MSE and MAE were 14 and 8.8 years, respectively. The performance of the NHP-E in situation (ii) is almost similar to that of the combined model. It can be concluded that when the failure intensity function is conditioned on the past failures, i.e. at least one failure event is included in the event sequence, the NHP model is able to make accurate predictions. The performance of the NHP-P is slightly better than its performance in situation (i). Furthermore, the NHP-P is still unable to distinguish whether to predict the first failure or a consecutive failure on the asset unit, based on its performance of a MSE 12.6 years and an MAE of 8.6 years.

The Neural Hawkes Process model in situation (ii) is able to predict the timestamp of the next failure based on the profile features and the past failures incorporated in the event sequence. In Figure 5, the time to next failure in years as observed, i.e. true time to next failure, and the time to next failure predicted by the Neural Hawkes Process model is shown. It can be seen that both distributions are different from each other. The actual time to next failure, shown in Figure 5a, indicates that most consecutive failures occur within 2 years with a maximum of 12 years. Figure 5b shows that the Neural Hawkes Process model predicts that most consecutive failures will occur between 2 and 4 years since the last failure, with

a maximum of 16 years.

C. Exciting and inhibiting effects of failure types

Lastly, the ability of the Neural Hawkes Process model to incorporate exciting and inhibiting effects of the failure event types on the failure intensity function is investigated. All failure events consist of a timestamp and a specific event type. In total seven failure types have been introduced in Figure 1. To determine the existence of these effects, the event type is excluded from the event sequence and the performance of NHP-C, is evaluated in both situations (i) and (ii).

The results are shown in Table II. Referring to situation (i), the effect of including the event types shows different results when the MSE and MAE are compared. First, the accuracy of the NHP-C, based on the MSE, decreases when the failure event types are included. However, the mean absolute error of the timestamp prediction improves both on the train and test set. In situation (ii), the performance accuracy has improved for both the MSE and MAE on the train and test sets.

Overall, it can be concluded that the performance of the event sequence is more leveraged in situation 2, wherein including the event types results in a better MSE and MAE score on both the train and test set. In this situation, the performance improvement, when the exciting and inhibiting effects of failure events are incorporated is clearly seen. In situation (ii), when the failure intensity function is conditioned on past failures, the performance on the model on the test set is 2.17 and 1.64 years, respectively the MSE and MAE. In case of situation (i), most event sequences have no actual failures during training. However, still the average error of the timestamp prediction is improved when the failure events are included, respectively 8.79 and 8.96 years.

TABLE II: Performance comparison on exciting and inhibiting effects

		Situation (i)		Situation (ii)	
		Including	Excluding	Including	Excluding
Train	MSE	13.55	13.76	2.15	2.17
	MAE	8.67	8.89	1.59	1.69
Test	MSE	13.99	13.82	2.17	2.26
	MAE	8.79	8.96	1.64	1.78

VI. CONCLUSION

We have shown that embedding the Hawkes process into a neural network structure is capable of modelling the failure intensity function of water pipes. The base failure intensity rate is learned based on a static profile of structural, environmental and operational water pipe attributes. The long-term dependencies of past failures on the failure intensity function is captured in a failure event sequence and feed into an LSTM. The performance of the proposed Neural Hawkes Process model has been evaluated using real data provided by a major Dutch water distribution network company. The results show

that the proposed model has achieved a better performance on predicting the time to the next failure.

As future work, more time-varying features such as temperature and climate can be incorporated into the RNN model, to increase the expressiveness of the model and to allow for a dynamic base failure rate.

REFERENCES

- [1] A. M. St. Clair and S. Sinha, "Development of a standard data structure for predicting the remaining physical life and consequence of failure of water pipes," *Journal of Performance of Constructed Facilities*, vol. 28, no. 1, pp. 191–203, 2014.
- [2] D. Winkler, M. Haltmeier, M. Kleidorfer, W. Rauch, and F. Tschekner-Gratl, "Pipe failure modelling for water distribution networks using boosted decision trees," *Structure and Infrastructure Engineering*, vol. 14, no. 10, pp. 1402–1411, 2018.
- [3] C. Zhang, H. Wu, R. Bie, R. Mehmood, and A. Kos, "Dynamic modeling of failure events in preventative pipe maintenance," *IEEE Access*, vol. 6, pp. 12 539–12 550, 2018.
- [4] P. D. Rogers and N. S. Grigg, "Failure assessment model to prioritize pipe replacement in water utility asset management," in *Water Distribution Systems Analysis Symposium 2006*, 2008, pp. 1–17.
- [5] A. Kumar, S. A. A. Rizvi, B. Brooks, R. A. Vanderveld, K. H. Wilson, C. Kenney, S. Edelstein, A. Finch, A. Maxwell, J. Zuckerbraun *et al.*, "Using machine learning to assess the risk of and prevent water main breaks," in *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. ACM, 2018, pp. 472–480.
- [6] S. Ertekin, C. Rudin, T. H. McCormick *et al.*, "Reactive point processes: A new approach to predicting power failures in underground electrical systems," *The Annals of Applied Statistics*, vol. 9, no. 1, pp. 122–144, 2015.
- [7] J. Yan, Y. Wang, K. Zhou, J. Huang, C. Tian, H. Zha, and W. Dong, "Towards effective prioritizing water pipe replacement and rehabilitation," in *Twenty-Third International Joint Conference on Artificial Intelligence*, 2013.
- [8] S. Xiao, J. Yan, X. Yang, H. Zha, and S. M. Chu, "Modeling the intensity function of point process via recurrent neural networks," in *Thirty-First AAAI Conference on Artificial Intelligence*, 2017.
- [9] H. Mei and J. M. Eisner, "The neural hawkes process: A neurally self-modulating multivariate point process," in *Advances in Neural Information Processing Systems*, 2017, pp. 6754–6764.
- [10] T. Xia, Y. Song, Y. Zheng, E. Pan, and L. Xi, "An ensemble framework based on convolutional bi-directional lstm with multiple time windows for remaining useful life estimation," *Computers in Industry*, vol. 115, p. 103182, 2020.
- [11] P. R. d. O. da Costa, A. Akcay, Y. Zhang, and U. Kaymak, "Remaining useful lifetime prediction via deep domain adaptation," *Reliability Engineering & System Safety*, vol. 195, p. 106682, 2020.
- [12] A. Shirzad, M. Tabesh, and R. Farmani, "A comparison between performance of support vector regression and artificial neural network in prediction of pipe burst rate in water distribution networks," *KSCE Journal of Civil Engineering*, vol. 18, no. 4, pp. 941–948, 2014.
- [13] D. Tran, A. Ng, B. Perera, S. Burn, and P. Davis, "Application of probabilistic neural networks in modelling structural deterioration of stormwater pipes," *Urban Water Journal*, vol. 3, no. 3, pp. 175–184, 2006.
- [14] B. Snider, E. A. McBean *et al.*, "Improving time to failure predictions for water distribution systems using extreme gradient boosting algorithm," in *WDSA/CCWI Joint Conference Proceedings*, vol. 1, 2018.
- [15] R. Jafar, I. Shahrour, and I. Juran, "Application of artificial neural networks (ann) to model the failure of urban water mains," *Mathematical and Computer Modelling*, vol. 51, no. 9–10, pp. 1170–1180, 2010.
- [16] Á. Martínez-Codina, M. Castillo, D. González-Zeas, and L. Garrote, "Pressure as a predictor of occurrence of pipe breaks in water distribution networks," *Urban Water Journal*, vol. 13, no. 7, pp. 676–686, 2016.
- [17] A. G. Hawkes, "Point spectra of some mutually exciting point processes," *Journal of the Royal Statistical Society: Series B (Methodological)*, vol. 33, no. 3, pp. 438–443, 1971.
- [18] P. NLextract. (2019) Basisregistratie adressen en gebouwen (bag) (uigebreed). [Online]. Available: <https://nlxtract.nl/downloads/>