Proceedings of the Twenty-Ninth AAAI Conference on Innovative Applications (IAAI-17))

Cracks Under Pressure? Burst Prediction in Water Networks Using Dynamic Metrics

Gollakota Kaushik*, **Abinaya Manimaran***
Arunchandar Vasan, Venkatesh Sarangan
TCS Research, IIT-Madras Research Park,
Chennai 600113, India
arun.vasan@tcs.com

Anand Sivasubramaniam
Dept. of Comp. Sci. & Eng.
Pennsylvania State University
University Park, PA 16802, USA
anand@cse.psu.edu

Abstract

Ranking pipes according to their burst likelihood can help a water utility triage its proactive maintenance budget effectively. In the research literature, data-driven approaches have been used recently to predict pipe bursts. Such approaches make use of static features of the individual pipes such as diameter, length, and material to estimate burst likelihood for the next year by learning over past historical data. The burst likelihood of a pipe also depends on dynamic features such as its pressure and flow. Existing works ignore dynamic features because the features need to be measured or are difficult to obtain accurately using a well-calibrated hydraulic model.

We complement prior data-driven approaches by proposing a methodology to approximately estimate the dynamic features of individual pipes from readily available network structure and other data. We study the error introduced by our approximation on an academic benchmark water network with ground truth. Using a real-world pipe burst dataset obtained from a European water utility for multiple years, we show that our approximate dynamic features improve the ability of machine learning classifiers to predict pipe bursts. The performance (as measured by the percentage of future bursts predicted) of the best forming classifier improves by nearly 50% through these dynamic features.

1 Introduction

Pipe bursts are a major problem for water utility networks. Water utilities address bursts either reactively post the event and fix the problem; or pro-actively by repair and rehabilitation of pipes over a much longer timescale, e.g., annually. The latter approach is preferred to avoid significant water loss and any penalties for damage to customer properties. Financially strong utilities can afford to pro-actively maintain a good part of the network. As an example, Sydney water, with an annual profit of AUD 513 million in 2014-2015 and a pipeline network of 21000 kms, inspected around 12034 kms over 2014-2015 (Sydney Water 2015). Such budgets are however outliers and specific to utilities in developed geographies that are water starved. The more common case is that a utility has a budget constraint (typically annual) and would like to thus spend the budget in minimizing the number of burst events over the next year.

Current Practice: To prevent bursts, we need to identify the pipes that have the highest risk (probability) of bursting the next year. Utilities typically have a deterministic set of rules that evaluate the risk associated with each pipe and evaluate a risk-reward matrix, for example (Anthony Cox 2008).

Researchers have proposed approaches for burst prediction that improve upon the risk-reward matrix. These approaches can be summarized as follows (more details follow in the next section). Low-level physically-based approaches such as (Kim et al. 2007; Rezaei, Ryan, and Stoianov 2015), model the pipe degradation by material properties. However, these approaches are more suited for better modeling of individual pipes in controlled environments and may not be scalable (and perhaps impractical) in real world settings.

Machine learning based approaches such as (Francis, Guikema, and Henneman 2014; Lin et al. 2015) model the pipes by learning the probability of failures from past history of bursts. The typical features include *static* pipe-level metrics such as length, diameter, age, etc., that are independent of other pipes.

Dynamic features: While data-driven methods have improved burst prediction over current industry practice, they ignore *dynamic* features of individual pipes such as pressure and flow that arise due to the interplay between the network structure and operations. The pressures and flows that a pipe experiences affects the amount of stress a pipe is subjected to and hence may play a role in its failure. Because pressure and flow of individual pipes arise from network-wide supply and demand constraints, they need to be either measured directly or computationally derived by solving a calibrated hydraulic model that simulates the operation of the network.

Challenges: Pressure and flow data is available only at the inlets of District Metering Areas (DMAs) (i.e., subnetworks) of the utility network. This is because pressure and flow sensors are typically deployed only at the DMA inlets (Narayanan et al. 2014). In many utilities including the one we study, the hydraulic model is typically not available for the entire network or is out of date due to the expensive calibration requirements. This could also be a reason why existing works do not take dynamic features into account.

In our work, we address the following problem: "Given historical burst data, basic pipe-level features, and the water

*The first two authors contributed equally to this work.

network structure, can we better estimate the probability of a pipe bursting in the next year using dynamic features even in the absence of a calibrated hydraulic model?" We complement prior machine learning based pipe classification approaches (burst or no-burst) by including dynamic pipe-level features such as pressure and flow that differentiate risk between pipes with identical static features being used differently in the network. As in existing approaches, the continuous output of the (burst or no-burst) classifier computed over a pipe's static and dynamic features is treated as the probability of that pipe's burst over the next year. Pipes are ranked according to the burst probability for any repair.

Contributions: To overcome the lack of a calibrated hydraulic model, we approximate the pressure and flow in each pipe of the water network using readily available data. Our intuition is that approximate values of the dynamic metrics that preserve the *relative ordering* across pipes should suffice for burst prediction. Using a benchmark water network where the ground truth of pipe flows and pressures are available, we demonstrate that our method of approximating the values of pressure and flow largely preserves the ordering of pipes by flows and pressures.

We evaluate the efficacy of these approximate dynamic features for predicting pipe bursts using a real-world data-set obtained from a leading European water utility. The data-set contains historical burst logs, static pipe features, and the network structure. After obtaining approximate pressures and flows, we train over a subset of the data and test over the remaining. We have tried approaches including both generative methods (Naive Bayes) and discriminative classifiers (Logistic Regression). The pipes are then ranked and assumed to be examined by the utility through a simulation in order of their ranking. Because pipes need to be tracked manually for any degradation, the cost of examining a pipe is proportional to the length of the pipe. Therefore, our performance metric for burst prediction is a modified performance curve that looks at fraction of bursts avoided per fraction of network length examined. Our **key findings** include the following:

- Approximate dynamic features consistently improves the accuracy of burst prediction for Naive Bayes across inspection lengths of interest to our water utility, *v*iz. $\leq 10\%$ of the total pipe length in the network;
- Naive Bayes with dynamic features shows the best performance. Specifically, nearly 30% of potential pipe breaks can be avoided by inspecting 10% of the pipe lengths. The use of dynamic features improves the performance of Naive Bayes classifier by 50% at this inspection length.
- Additional pipes that are correctly identified by our method are characterized by being in the upper tail of the distributions of the approximate dynamic features, confirming the usefulness of the dynamic features.

Significance. Water utilities operate on a shoe-string budget as the economic value of water is not as high as energy. Any improvement in their ability to correctly identify the pipes in need of proactive maintenance will greatly increase their operational efficiency. We find that the approximate dynamic features allow a utility to accurately identify additional potential pipe bursts for smaller network inspection lengths. We note here that in case a utility can accurately estimate the dynamic features through well calibrated hydraulic models, the usage of such accurate dynamic feature values will only further improve the prediction accuracy. Based on the proposed methodology, we are developing a software application for possible deployment with a real world water utility for proactive pipe maintenance. We also plan to enhance the software application by experimenting with more sophisticated approaches such as deep learning and boosting to further improve the burst prediction performance.

2 Related work

Physically-based models: Physically-based models estimate the wear and tear of a pipe based on the mechanical strength of pipe materials; and the usage of the pipe. Finite element 3-D modeling is used in (Robert et al. 2016) for predicting pipe stress and thus pipe bursts after accounting for internal and external loading of pipes. Reference (Davis et al. 2007) presents a physical probabilistic model using fracture mechanics theory. Growth of cracks using visual micro-scale examination is used in (Gould et al. 2013) to do failure analysis of PVC sewer lines. The correlation between dynamic pressure in the network and the crack development mechanisms is studied in (Rezaei, Ryan, and Stoianov 2015). The residual tensile strength of iron pipes is studied in (Kim et al. 2007) to predict bursts; A comprehensive survey of physically-based models for pipe deterioration is presented in (Rajani and Kleiner 2001). While physical approaches can be accurate for specific pipes under study, it is unclear if they are scalable across pipes of multiple types, lengths, ages, and diameters in a typical network. This is because, real world networks may have thousands of pipes with different materials, age, size, length, stress, and soil conditions. It may not be possible to develop failure individual models for each of these pipes. Moreover, measuring the parameters required for such physical models for individual pipes under operation may not be possible.

Machine learning & Statistical models: One of the earliest works is (Kettler and Goulter 1985) which models the probability of bursts using regression on the pipe features. Reference (Kleiner and Rajani 1999) shows how to use a limited data approach to modeling bursts in homogeneous pipe groups and fits exponential failure rate models to pipes. In (Lin et al. 2015), the authors use a Dirichlet process mixture of hierarchical beta processes to estimate the burst likelihood. Additional features like soil type and proportional hazard models are used in (Yamijala, Guikema, and Brumbelow 2009). Multiple methods for classification and rank boosting are tried in (Wang et al. 2013); the authors report that boosting significantly improves the performance. Unlike other approaches which look for temporal patterns, (de Oliveira et al. 2010) looks for spatial patterns across the distribution network. Belief networks are used in (Francis, Guikema, and Henneman 2014) for prediction of pipe breaks. Artificial neural networks have been used in (Harvey, McBean, and Gharabaghi 2013) to predict mains fail-

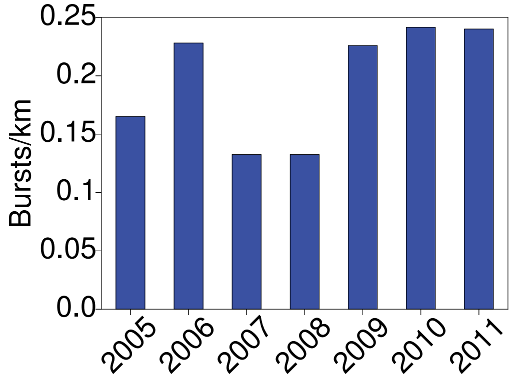

Figure 1: Bursts per km of pipe length in each year. The data is shown for large diameter pipes.

Symbol	Meaning
E	Pipes (edges, links)
V	Junctions (vertices, nodes)
d_i	Demand at node i
p_i	Static pressure at node i
h_i	Elevation at node i
ρ	Density of water (1 gm/cc)
g	Acceleration due to gravity 9.8 m/s^2
H_i	Total pressure at node i equals $p_i + h_i\rho g$
$C_{i,j}$	Pressure gain due to pump along pipe (i, j)
$f_{i,j}$	Flow from node i to node j
e	Edge (pipe) $e = (i, j)$
L_e	Length of pipe e
R_e	Radius of pipe e
μ_e	Roughness coefficient of pipe e

Table 1: Notation used in the paper.

ures. Reference (Yan et al. 2013) shows that machine learning based classifiers better predict pipe bursts than conventional risk-matrix based approaches used by the utilities.

A common limitation of these data-driven approaches is that they focus on static metrics that are pipe specific without considering dynamic metrics.

3 Description of real-world data-set

The real world pipe burst data-set obtained was from an European water utility[1]. The network has around 6785 pipes spanning 415 kilometers of length. The network is divided into 16 DMA's (District Metering Areas) which are subnetworks over which flows can be controlled and monitored at few points of entry. For each DMA, the average flow entering the DMA (over a typical day) was known. The static data included the following: length (0.5-1580 meters), diameter (13-600mm), depth of burial (0.6-1.5 m), material (Cast Iron, PVC, etc), and the approximate date of laying the pipes (1930-2011).

A GIS system provided the connectivity information between the various pipes and identified the inlets to each of the DMA's. The rehabilitation planning is done by the utility and approved by the regulator annually, so we predict bursts one year ahead. Burst data logged digitally was available from 2005.

Bursts in large diameter pipe (i.e., diameter $>$ 125mm) create more damage than a burst from a small diameter pipe. In terms of impact, there is a 80-20 rule in favor of the large diameter pipes – i.e., though the larger diameter pipes may account for smaller percentage of the overall bursts, they contribute to a major fraction of the resulting impact. Large diameter pipes are also more expensive to fix (The Water Research Foundation 2016). Given these, in the rest of this paper, we focus on predicting bursts in large diameter pipes. This focus on larger diameter pipes is similar to the approach followed in (Lin et al. 2015). There were around 175 bursts in large diameter pipes spread over 2005-2011. The number of bursts per kilometer of pipe length across large diameter pipes is shown in Figure 1. We observe that the probability of a pipe bursting (normalized with respect to length) is less when compared to probability of a pipe not bursting. Clearly, the priors of the two classes are imbalanced in favor of the non-burst class.

Finally, we obtained the elevation data for the entire network from Google maps at the spatial resolution of 150 meters. to model the effect of elevation on flow/pressure in the pipes.

4 Fundamentals of water networks

Water flows in the network under pressure (due to elevation or pumping). Each demand point on the network acts as a sink to the flows (and returns the flow to sewage network) and the demands at various points together influence the flows direction and magnitudes in the network. Each pipe offers a frictional resistance to the flow and thus reduces the pressure along the flow. The resistance would depend on the roughness coefficient (explained below), length, and diameter of the pipe. The demands at all nodes, the reservoirs' capacities and pressures, the rating of all pumps in a water network, and the pipe characteristics including the roughness coefficients are typically *the known or assumed inputs of a water network's hydraulic model.* The pressures at nodes and the flows in pipes are outputs to be determined from these inputs. Using the notation in Table 1, the equations relating the outputs with the inputs are (Prabhata Swamee and Ashok Sharma 2008):

- *Flow equation*: The flow entering a node i from all its neighbors $\mathcal{N}_i$ is equal to the total flow leaving it plus the demand d_i of the node i. In other words,

$$\forall i \in V, \ \ d_i + \sum_{j \in \mathcal{N}_i} f(i, j) = 0 \tag{1}$$

  This assumes that there are no leaks at junctions.

- *Pressure equation*: For a link $e = (i, j) \in E$, the change in the total pressure H along the link is given by:

$$H_i + PumpingGain - Friction = H_j \tag{2}$$

[1] For confidentiality reasons, the details are anonymized.

With the Total Pressure H_i re-written as the sum of the the static pressure p_i and the pressure due to elevation $h_i\rho g$, the previous equation is expanded as:

$$p_i + h_i\rho g + C_{i,j} - \frac{L_e f_e^2 \times \rho g}{\mu_e R_e^5} = p_j + h_j\rho g \quad (3)$$

where μ_e represents the pipes roughness coefficient.

These equations can be solved using a fixed point iteration – for example Todini's approach (Todini and Pilati 1987), to obtain the pressures at all nodes; and flows through all pipes.

Model calibration challenges: Equations 1 and 3, require the roughness coefficients μ_e's and the demands d_i's. These d_i's and μ_e's are empirically estimated during the process of hydraulic model calibration. In most water networks, customers are not precisely metered at end-points and even if they are, the data will typically be a monthly average. Demands are accurately measured only at aggregated levels. Therefore, node-level demands (d_i's) may not be accurately known at the required spatio-temporal granularity. Further, as pipes age and corrode, their roughness needs to be periodically estimated. Considerable effort is required to estimate the μ_e's. For example, calibrating a network with few hundreds of nodes may take 40-60 days of effort by an expert team of 2-4 members (Bros and Kalungi 2010). A utility lacking in-house expertise would require expensive niche consultants. For a network serving a population of 1 million, hydraulic model maintenance over a five year period typically costs about $4 million (Narayanan et al. 2014). Because of such reasons, all water utilities may not always have a well calibrated hydraulic model to accurately estimate the dynamic features.

5 Approximation of dynamic feature values

The data set available with the case study real-world water utility (as described earlier) is quite basic – network structure, pipe details, burst records, and readings of flow meters at DMA inlets. No calibrated hydraulic model was available. We now present an approach to approximately estimate the values of dynamic features (pressure[2] and flow) for each pipe even in the absence of an accurate hydraulic model.

The network structure from the GIS system is a logical graph, where the pipes are edges; and intersections of pipes are nodes. We create a pseudo-hydraulic model with the following simplifying assumptions: 1) Demands are uniformly distributed along the pipes of the lowest diameters. Thus the average inflow into a DMA is divided across all demand points. and 2) The roughness coefficients of the pipes have not changed significantly from the initial values at the time of deployment (which can known from the material type).

With these simplifying assumptions, we solve for an approximate hydraulic model using fixed point iteration and obtain the pressures at nodes and flows through various pipes. Note that if demands are varying with time[3] we could evaluate the time-averaged values of the pressures and flows through the pipe network.

Validation: To evaluate our approximation, we now check if these approximate pressures and flows are significantly different from reality. For this evaluation, we use a benchmark water network with a well calibrated hydraulic model from the University of Exeter(Center for Water Systems, University of Exeter). The Wolf-Cordera (Colorado Springs) network has 1981 pipes delivering 3.7 million gallons per day. Using the calibrated model, we obtain the accurate pressures and flows across the pipes – these actual values act as the ground truth. For this network, we also obtain the approximate flows and pressures under our two simplifying assumptions – these values act as our approximations for the ground truth.

Figures 2(a) through 2(e) compare the actual and approximate values of the dynamic features as a X-Y scatter plot. The dynamic features we consider are: (a) flow through a pipe, (b) total pressure for a pipe (midpoint), (c) static pressure for a pipe (midpoint), (d) total pressure drop across a pipe, and (e) static pressure drop across a pipe. The accurate values are in the X axis and the approximate values are in the Y axis. A reference $45°$ line is also shown in the plots to highlight the extent to which the approximate values deviate from the actual values. Larger the deviation, farther will be the points in the plot from the $45°$ line. As we can see from the plots, across all dynamic features, almost all the points lie along the $45°$ line. This indicates that the approximate values match reasonably well with the actual values.

We also further check if the relative ordering between the pipes is maintained when using the approximate features. For each dynamic feature F, we first pick the set of top k pipes (ranked by F) in both the approximate model $\mathcal{E}_k$ and the ground truth $\mathcal{A}_k$. Then we compute the *extent of intersection* $\frac{|\mathcal{E}_k \cap \mathcal{A}_k|}{k}$ as a measure of the usefulness of our approximation in identifying the top k pipes according to feature F. Figure 3 shows this measure for various dynamic features. Each curve shows the result for a specific feature.

Apart from initial values of k, the overlap between the sets identified by the approximate and actual models is close to around 95%. Even for high values of k, the minimum overlap between the sets is around 90%. The ordering is such that some pipes are initially omitted and then later added to the number of pipes being considered. So there is a dip and then the percentage comes back up. In sum, we note that despite the errors in estimating demands and the roughness coefficients the *relative ordering across pipes is preserved* according to the ranking for various dynamic features.

We hypothesize that standard machine learning classifiers can use these approximate dynamic feature values in place of actual feature values to predict future pipe bursts. We evaluate this hypothesis in the next section.

6 Burst prediction using dynamic features

We now evaluate the use of approximate dynamic features to predict bursts. The real world dataset spans from 2005 to 2010. The water utility has to decide the maintenance for the next year (e.g., 2010) at the end of current year (e.g., 2009).

[2] By pressure, we mean the pressure at the midpoint of the pipe.

[3] In our dataset, we have only the average inflow into the DMAs.

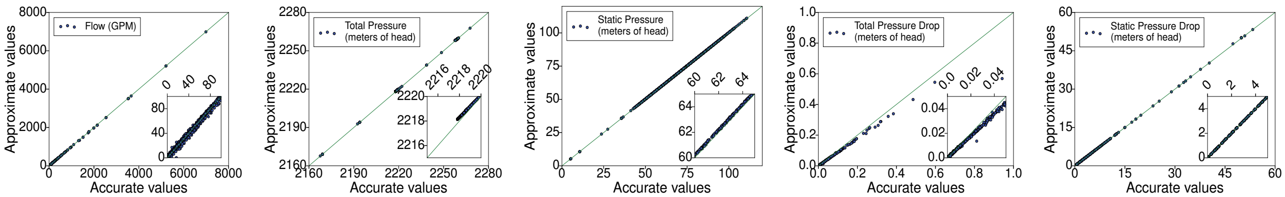

(a) Flow (b) Total pressure (c) Static pressure (d) Total pressure drop (e) Static pressure drop

Figure 2: Scatter plot between the approximate and actual values of the dynamic features. The scatter lies close to the $45°$ line indicating that the approximate values reasonably match with the actual feature values.

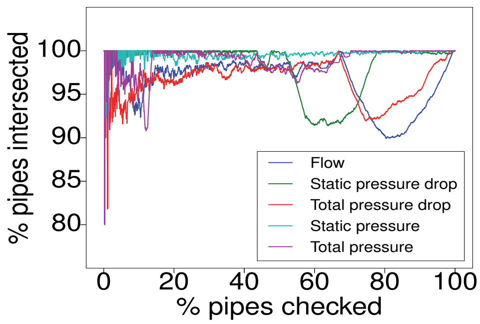

Figure 3: Ordering is preserved across pipes despite approximations in the dynamic feature values.

Static Features	**Value type**
Length	Continuous
Diameter	Continuous
Age	Continuous
Connectivity	Continuous
Material	Categorical
Depth	Continuous
# of bursts occurred so far	Continuous
Dynamic Features	**Value type**
Total Pressure	Continuous
Static pressure	Continuous
Total pressure drop	Continuous
Static pressure drop	Continuous
Flow	Continuous

Table 2: Features available for burst prediction.

Hence the prediction will use the available data/features of all the pipes laid until the end of the current year (2009) and predict their ranks based on the break probabilities for the next year (2010). Correspondingly, the training set for the model would consist of feature values of all pipes that were laid until the end of last year (2008) and their break/class information from the current year (2009).

Data set correction: The utility under study already has a proactive maintenance program wherein high risk pipes are identified as per a 'risk-reward' matrix. Among the identified pipes, the top ranking pipes totaling $\approx 1\%$ of the total network length are proactively inspected and fixed. Consequently, the burst data set provided by the utility is biased by the proactive pipe maintenance activities. In order to remove the effects of this bias, all the pipes that are proactively maintained by the utility in a year n are removed for both testing and training for years greater or equal to n.

Classifiers used: To evaluate the approximate dynamic features in predicting future pipe bursts, we tested them with four different classifiers – Naive Bayes, Restricted Boltzmann machines, Logistic Regression and Support Vector Machines. Due to lack of space, we discuss results from the classifier that performed the best in each category – one *generative, namely Naive Bayes* and one *discriminative, namely Logistic Regression*.

Feature Selection: For each pipe, we have a total of 7 static features – namely length, diameter, age, connectivity, material, depth, and number of bursts occurred so far. Using all features to predict future pipe bursts from historical data need not necessarily improve a classifier's performance.

Methods have been proposed in the literature to choose features for classification in a systematic way. In our case, however, the data-set is low-dimensional, e.g., 7 static features. Therefore, we bypass the problem of feature selection by considering *all* $2^7 - 1$ *combinations* of features and identifying the best-performing combination of base features for each classifier type over the training set. Similarly, we have a total of 5 dynamic features for each pipe – namely total pressure, static pressure, total pressure drop, static pressure drop and flow. With the additional dynamic features, we identify the best feature set combination (out of the $2^{12} - 1$ choices) that has at least one dynamic feature. The category of various static and dynamic features used is summarized in Table 2.

Quantifying classifier performance: The classifier predicts the class of each pipe (i.e., burst or not) with an output probability. The output of the classifier is used as the rank of the pipe that is classified. In all these approaches, we first train a classifier on past-data and then evaluate the classifier on a future year. Suppose k pipes of lengths $L_1 \dots L_k$ are examined first according to the ranking given by a classifier. These set of k pipes would entail certain percentage of actual bursts in the coming year (which is available from the ground truth data). Combining these two will result in a point on the performance curve. Specifically, the X-coordinate of the point is given by $\frac{\sum_{i=1}^{i=k} L_i}{\sum_{i=1}^{i=N} L_i}$, that is the frac-

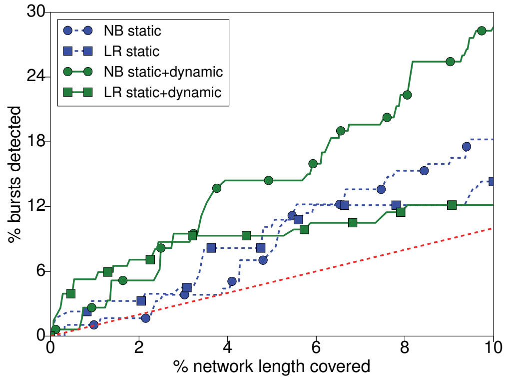

Figure 4: Performance curves for dynamic features under two different classifiers. The performance curves are obtained by averaging the classifiers' performance across the years 2005-2010.

tion of the total network length examined. The Y-coordinate shows the percentage of bursts in these k examined pipes as validated from the ground truth. Note that pipes that were fixed proactively by the utility are removed from consideration in both training and testing. This process is repeated for various ranks, $k = 1 \ldots N$, and the performance curve is obtained. To summarize, the X-axis in each of the figures is the % of the length of the network examined. The Y-axis shows the percentage of predicted bursts that are actually bursts.

7 Results

As with recent approaches (Yan et al. 2013), we too observed that predicting pipe bursts with machine learning based classifiers using static features outperform conventional risk-reward approaches. For the sake of brevity, in this work, we focus on comparing the performance of classifiers under two scenarios: (a) using static features, and (b) using both static and dynamic features.

Naive Bayes (NB): Figure 4 shows the average performance of NB with and without dynamic features across the years 2005-2010. The best performing dynamic feature combination included diameter, depth, material, connectivity, static pressure and static pressure drop. We find that the use of dynamic features (though approximate) consistently and considerably improves the performance of the classifier across the inspection lengths of interest to the utility ($< 10\%$). At an inspection length of 10%, the dynamic features allow the classifier to identify nearly 50% additional bursts on average in comparison to static features. NB uses static pressure and static pressure drop across a pipe as dynamic features to improve the classification performance. As a case in point, in the year 2009, when 8% of the network length is examined, NB detects 32% of pipe bursts using dynamic features, while only 8% of the burst pipes are detected using static features. Nearly two thirds of the bursts detected by the dynamic feature combination have dynamic feature values that lie in the upper tail of the feature distribution which improves the classifier's performance. Though there are some common break pipes that are detected by both dynamic and static feature combination, the ranks of those pipes are moved significantly to the top order when dynamic features are considered. This helps the utility to examine even a lesser percentage of the network.

Logistic Regression (LR): Figure 4 shows the average performance of the LR with and without dynamic features across the years 2005-2010. The best performing combination of features included the static feature depth; and the dynamic features flow and total pressure. We find that the use of dynamic features improves the performance when smaller fractions of the network are taken up for inspection by the utility ($< 5\%$). At an inspection length of 2%, the dynamic features allow LR to identify nearly 50% additional bursts on average in comparison to static features. As the inspection length increases, the use of dynamic features does not add much value. In other words, LR is not able to learn the information content in the dynamic features and combine it appropriately with those in the other static features. The dynamic features used by logistic regression to give an improved performance at smaller inspection lengths are the total mid-point pressure and the flow value. Specifically, during the year 2009, on examining 4% of the network length, logistic regression with dynamic features detected nearly 25% of the total bursts while using static features alone, the classifier should detect just 5% of the total bursts. This performance boost comes because nearly 66% of the breaks detected by using dynamic features occurred on pipes which belong to the top 10 percentile of the dynamic features.

Summary: We observe the following:

- Dynamic features do improve the performance of classifiers across the inspection lengths of interest to the utility – for smaller inspection lengths ($< 3\%$), logistic regression with dynamic features gives the best performance, while for higher lengths, Naive Bayes with dynamic features gives the best performance. This suggests that even *approximate network operational information can help utilities improve their burst prediction.*
- Depending on the inspection length desired by the utility, the appropriate classifier with the best performance can be chosen. If simplicity is desired, on account of consistent performance, Naive Bayes can be picked as the single classifier to do predict pipe bursts.
- While the performances of the two classifiers are comparable under static features, they significantly differ when approximate dynamic features are introduced. This indicates that different classifiers make use of the additional information available from the dynamic features differently across the spectrum of inspection lengths.

8 Conclusions

Water utilities need to prioritize rehabilitation of pipes for a given maintenance budget. Since they are budget constrained, any improvement in their ability to correctly identify the pipes in need of proactive maintenance will greatly increase their operational efficiency. To this end, we presented an approach that approximates the dynamic features of pipes, specifically flow and pressure, using readily available data. We showed that the performance of machine learning classifiers to predict future pipe bursts significantly improves when trained with our approximate dynamic features. We are in the process of deploying this technique in real world water utilities as a separate software application for proactive pipe maintenance. Future directions of work include experimenting with the use of approximate dynamic features in techniques like boosting and deep learning to further improve the prediction performance. We are also developing a software application based on the proposed methodology for possible deployment with a real world water utility for proactive pipe maintenance.

References

Anthony Cox, L. 2008. What's wrong with risk matrices? *Risk analysis* 28(2):497–512.

Bros, C. M., and Kalungi, P. 2010. Calibration and sensitivity analysis of the c-town pipe network model. *ASCE Water Distribution Systems Analysis 2010* 1507–1523.

Center for Water Systems, University of Exeter. Wolf-Cordera Ranch.

Davis, P.; Burn, S.; Moglia, M.; and Gould, S. 2007. A physical probabilistic model to predict failure rates in buried pvc pipelines. *Reliability Engineering & System Safety* 92(9):1258–1266.

de Oliveira, D. P.; Neill, D. B.; Garrett Jr, J. H.; and Soibelman, L. 2010. Detection of patterns in water distribution pipe breakage using spatial scan statistics for point events in a physical network. *Journal of Computing in Civil Engineering* 25(1):21–30.

Francis, R. A.; Guikema, S. D.; and Henneman, L. 2014. Bayesian belief networks for predicting drinking water distribution system pipe breaks. *Reliability Engineering & System Safety* 130:1–11.

Gould, S. J.; Davis, P.; Beale, D. J.; and Marlow, D. R. 2013. Failure analysis of a pvc sewer pipeline by fractography and materials characterization. *Engineering Failure Analysis* 34:41–50.

Harvey, R.; McBean, E. A.; and Gharabaghi, B. 2013. Predicting the timing of water main failure using artificial neural networks. *Journal of Water Resources Planning and Management* 140(4):425–434.

Kettler, A., and Goulter, I. 1985. An analysis of pipe breakage in urban water distribution networks. *Canadian Journal of Civil Engineering* 12(2):286–293.

Kim, J.; Bae, C.; Woo, H.; Kim, J.; and Hong, S. 2007. Assessment of residual tensile strength on cast iron pipes. In *ASCE Int. Conf. on Pipeline Eng. and Construction, Boston.*

Kleiner, Y., and Rajani, B. 1999. Using limited data to assess future needs. *American Water Works Association. Journal* 91(7):47.

Lin, P.; Zhang, B.; Wang, Y.; Li, Z.; Li, B.; Wang, Y.; and Chen, F. 2015. Data driven water pipe failure prediction: A bayesian nonparametric approach. In *Proc. of the 24th ACM Int. Conf. on Information and Knowledge Mgmt.*, 193–202.

Narayanan, I.; Vasan, A.; Sarangan, V.; and Sivasubramaniam, A. 2014. One meter to find them all: water network leak localization using a single flow meter. In *Proc. of the 13th Int. Symposium on Information Processing in Sensor Networks, April 15-17, 2014, Berlin*, 47–58.

Prabhata Swamee and Ashok Sharma. 2008. *Design of Water Supply Pipe Networks*. Wiley-Interscience.

Rajani, B., and Kleiner, Y. 2001. Comprehensive review of structural deterioration of water mains: physically based models. *Urban water* 3(3):151–164.

Rezaei, H.; Ryan, B.; and Stoianov, I. 2015. Pipe failure analysis and impact of dynamic hydraulic conditions in water supply networks. *Procedia Engineering* 119:253–262.

Robert, D.; Rajeev, P.; Kodikara, J.; and Rajani, B. 2016. An equation to predict maximum pipe stress incorporating internal and external loadings on buried pipes. *Canadian Geotechnical Journal* (ja).

Sydney Water. 2015. Water Efficiency Report.

The Water Research Foundation. 2016. Asset Management Breaks and Leaks Factsheet. http://www.waterrf.org/knowledge/asset-management/factsheets/assetmgt-breaksleaks-factsheet.pdf.

Todini, E., and Pilati, S. 1987. A gradient method for the analysis of pipe networks. In *International Conference on Computer Applications for Water Supply and Distribution, Leicester Polytechnic, UK*.

Wang, R.; Dong, W.; Wang, Y.; Tang, K.; and Yao, X. 2013. Pipe failure prediction: A data mining method. In *Data Engineering, 29th Int. Conf. on*, 1208–1218. IEEE.

Yamijala, S.; Guikema, S. D.; and Brumbelow, K. 2009. Statistical models for the analysis of water distribution system pipe break data. *Reliability Engineering & System Safety* 94(2):282–293.

Yan, J.; Wang, Y.; Zhou, K.; Huang, J.; Tian, C.; Zha, H.; and Dong, W. 2013. Towards effective prioritizing water pipe replacement and rehabilitation. In *IJCAI 2013, Proc. of the 23rd International Joint Conference on Artificial Intelligence, Beijing*, 2931–2937.