A Noise Scaled Semi Parametric Gaussian Process Model for Real Time Water Network Leak Detection in the Presence of Heteroscedasticity

Obaid Malik, Siddhartha Ghosh, Alex Rogers
University of Southampton
UK, SO17 1BJ
(om1e09, sg2, acr)@ecs.soton.ac.uk

Abstract

The timely detection of leaks in water distribution systems is critical to the sustainable provision of clean water to consumers. Increasingly, water companies are deploying remote sensors to measure water flow in real-time in order to detect such leaks. However, in practice, for typical District Metering Zones (DMZ), financial constraints limit the number of deployable real time flow sensors/meters to one or two, thus constraining leak detection to be based on the aggregated flow being monitored at these point. Such aggregated flow data typically exhibits input signal dependence whereby both noise and leaks are dependent on the flow being measured. This limited monitoring and input signal dependance make conventional approaches based on simple thresholds unreliable for real time leak detection. To address this, we propose a Gaussian process (GP) model with an additive diagonal noise covariance that is able to handle the input dependant noise observed in this setting. A parameterised mean step change function is used to detect leaks and to estimate their size. Using prior water distribution systems (WDS) knowledge we dynamically bound and discretize the detection parameters of the step change mean function, reducing and pruning the parameter search space considerably. We evaluate the proposed noise scaled GP (NSGP) against both the latest research work on GP based fault detection methods and the current state of the art and applied leak detection approaches in water distribution systems. We show that our proposed method out performs other approaches, on real water network data with synthetically generated time varying leaks, with a detection accuracy of 99%, almost zero false positive detections and the lowest root mean squared error in leak magnitude estimation (0.065 l/s).

1 Introduction

The provision of clean water has been a primary source of concern for civilisations since ancient times. Despite major advances in science and technology over the centuries, many parts of the world still suffer from clean water shortages. The World Health Organisation estimates that 783 million people (i.e. one in ten) in the world are deprived of safe water. In light of this scarcity, the efficient use and conservation of water resources is of utmost importance. Leaks are a major cause of water loss. In urban areas leakage in water distribution networks can be as high as 70% (SWAN 2011), thus early detection of leaks is critical to maintaining water supplies. In the context of water utility companies, leaks result in revenue loss. An obvious solution is to install sensors at every location within a District Metering Zones (DMZ) - an isolated monitoring area consisting of a few 100 to 2000 properties, which would facilitate the rapid detection of leaks at street level. However, this is often financially infeasible. Therefore typically, just one flow meter is installed at the point where the DMZ connects to the rest of the network. The aggregated flow readings at this point represent the current water demand/consumption within the DMZ. Detecting leaks at this point is challenging, particularly small leaks, as they do not produce a significant increase in the aggregated DMZ level flow. Also, because flow and pressure within the DMZ are not only correlated but also affect the noise levels associated with the readings being taken, such that - an increase in the demand/flow decreases pressure and an increase in the supplied pressure increases the flow within the pipes, in turn not only increasing the noise levels associated with the readings being taken but also the leak magnitude (if a leak is present in the system at that time), both leaks and the noise in the observed flow readings in the system are dependent on the inputs. This input dependent and time varying noise is often referred to as heteroscedasticity). Besides leaks, genuine customer consumption can also cause an unexpected increase in the flow e.g. fire fighting, or high demand during a festival. Thus the duration of the unexpected increase must be taken into account before attributing it to a leak. For financial accountability, quantification of the water loss during a leak, is of equal importance for the water companies. Given the nature of the leaks and monitoring requirements in a DMZ we can list the requirements for a practical leak detection solution:

1. Timely detection of leaks when only limited water network data is available i.e using only aggregated DMZ level flow/pressure measurements.
2. Distinction of legitimate short term increase in water consumption from leaks.
3. Estimation of the leak magnitude to quantify water loss.

Now, there is a large collection of research work from various domains applied to, both the general area of fault/anomaly detection and the specialised area of leak detection in water networks (Zhang and Ding 2008; Chandola, Banerjee, and Kumar 2009; Hwang et al. 2010; Verde et al. 2008; Colombo, Lee, and Karney 2009; Puust et al. 2010; De Silva et al. 2011). However, these approaches are unsuitable and generally perform poorly as suffer from one or more of the following shortcomings:

1. Gaussian noise assumptions: Given the heteroscedasticity of this setting, detection approaches that assume the observation noise to be a constant Gaussian result in false positive detections at times of high variations in the consumption (water flow rate).
2. Assumption of leak independence: Most approaches in the literature assume the leak at time t to be independent of leaks at other times. Thus, they perform a point estimate of the leak at every new observed flow reading independent of the historic data. However, leaks in reality remain in the system until they are fixed. Given our requirements, point estimation of leaks can result in legitimate short term increase in consumption being flagged as a leak (false positives) or non-detection for small leaks at times of high flow variance (false negatives).
3. Timely/online leak detection: Leaks must be detected and reported in realtime. Some approaches only use a subset of data (e.g. one minimum variance point per day) to detect leaks or cannot be run in an online setting (batch mode detection approaches). These approaches normally result in unacceptable large leak detection/reporting times e.g. 24 hour or more.

Now, leak detection from flow data is a form of time series analysis and in this space Gaussian processes (GP) are gaining much attention, as they facilitate flexible and analytic inference in a fully Bayesian nonparametric setting. Recent works on fault/change point detection using GPs, have tried to address problems similar to leak detection (Garnett et al. 2010; Osborne, Garnett, and Roberts 2010; Osborne et al. 2012). However, these approaches, when applied to leak detection, again suffer due to the Gaussian noise and leak independence assumptions discussed above.

To address these shortcomings, in this paper, we propose an alternate GP based model which exploits the fact that water consumption (within a DMZ) follows a deterministic pattern which is periodic at seasonal, monthly and weekly scales. To incorporate this, we firstly use a weekly moving window on the data being modelled. Secondly we define two vectors, the weekly mean and weekly variance which are computed from the historic aggregate flow data. We then use the weekly mean vector as a GP prior mean. This methodology intrinsically incorporates seasonal and monthly variations in to the GP. Now, in a standard GP, the observation noise is assumed to have a constant Gaussian variance. Modelling heteroscedastic noise in a GP results in a non-closed form likelihood expression and intractable integrals for the posterior. To address this approximate methods to incorporate heteroscedastic noise have been proposed (Goldberg, Williams, and Bishop 1997; Kersting et al. 2007; Lazaro-Gredilla and Titsias 2011). However, it has been shown that these approximations and/or variational based approaches in GP result in under estimation of the posterior variance, which in case of leak detection would result in false positives (Kuss and Rasmussen 2005; Consonni and Marin 2007). To avoid this problem, we model the noise variance in the GP as an additive diagonal noise covariance kernel (Venanzi, Rogers, and Jennings 2013). This not only allows us to model the input dependent noise more accurately but also results in a closed form likelihood expression and tractable inference. Subsequently, leak detection is done via an additive parameterized step change mean function. The parameterization captures the properties of the leak we are interested in learning e.g. leak start time and leak magnitude, which allows us to not only detect a leak, but also, quantify the leak magnitude. The resulting model has the advantage of modelling a leak as being independent of the underlying normal operational model. This leads to a more accurate depiction of the physical leak process, since leaks result in sudden sustained additions uncorrelated to the original underlying consumption patterns. For each new observed flow reading, based on our knowledge of the WDS, we place dynamic bounds and then discretize the leak magnitude parameter (details in Section 5). This allows us to automatically prune our leak parameters search space resulting in a robust and fast detection approach that can be used in an online setting. Given this, we extend the state of the art in the following way

- We propose a novel GP based model, the Noise Scaled GP (NSGP), for water distribution systems that can handle heteroscedasticity and has a closed form likelihood expression resulting in tractable inference.
- By parameterizing the proposed GP, we model the physical leak process in a way that allows us to use WDS domain knowledge to dynamically bound the leak detection parameters, at each step, resulting in a fast leak detection and quantification method that can manage sustained leaks with high detection accuracy.
- By comparing the performance of our NSGP with both state of the art in GP based fault detection algorithms and current applied leak detection approaches in the water industry, on real data with simulated time varying leaks, we show that our proposed NSGP outperforms the state-of-the-art by achieving detection accuracies of 99%.

The remainder of the paper is arranged as follows. In Section 2 we introduce GP and briefly discuss hyper-parameter learning. Section 3 describes our Noise Scaled GP model for modelling water network data. In Section 4 we detail how the proposed model can be parameterized to detect leaks. Section 5 describes the hyper parameter learning schemes we use. Our results and comparison with selected approaches is presented in Section 6.

2 Gaussian Processes

Formally a GP can be defined as a stochastic process defining a distribution over functions $H \to \mathbb{R}$ such that the GP evaluated at any finite subset $F \subset H$ is a multivariate Gaussian distribution. A GP can be completely defined by a mean function, $m(.)$ and a positive semi-definite covariance function, $k(.,.)$. Given the observed values, $y = y_1 \dots y_n$, of a function f at a set on inputs $x_1 \dots x_n$, the observed sample

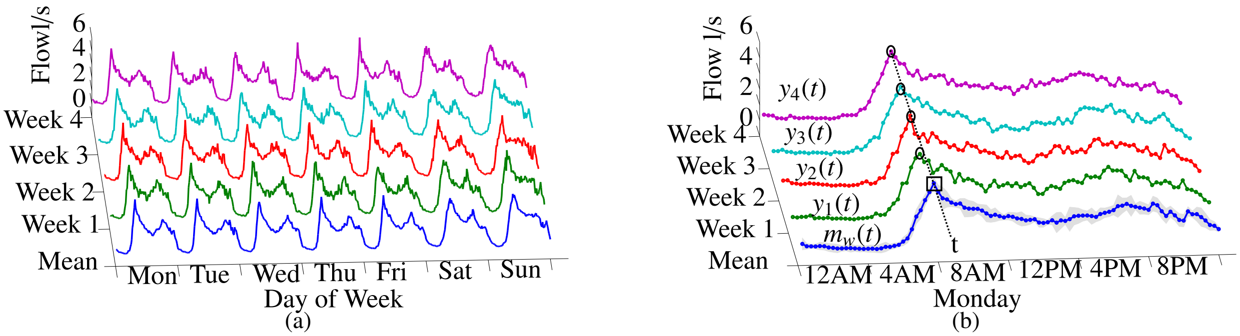

Figure 1: (a) Four week flow data showing similarity in weekly consumption patterns and the calculated weekly mean. Flow measured in litres per second (l/s). (b) Showing consumption patterns on Monday across four weeks and the computation of weekly mean and variance vectors.

can be thought as being drawn from a multivariate Gaussian distribution. Assuming a GP prior on f the prior distribution can be given as:

$$p(f|x,\theta) = \mathcal{N}\left(m(x;\theta), k(x,x;\theta)\right)$$

where $k(x,x;\theta)$ is the positive semi-definite covariance matrix and θ is the set of hyper-parameters that characterise the mean and covariance functions. However, often the observed data, y, is a noise corrupted version of the true underlying function f which is not known. In such cases, if the noise is assumed to be zero mean Gaussian i.i.d, $\epsilon \sim \mathcal{N}(0,\sigma^2)$, then the prior distribution can be written as:

$$p(y|x,\theta) = \mathcal{N}\left(m(x;\theta), k(x,x;\theta) + \sigma^2 I\right) \tag{1}$$

It can be shown (Rasmussen and Williams 2005) that the posterior predictive density given a set of test points x_* can be given as:

$$p(y_*|x_*,x,y,\theta) = \mathcal{N}\left(y_*; m(x_*|y,x,\theta), \Sigma(x_*|y,x,\theta)\right) \tag{2}$$

$$m(x_*|y,x,\theta) = m(x_*) + K(x_*,x)K_y^{-1}(y - m(x)) \tag{3}$$

$$\Sigma(x_*|y,x,\theta) = K(x_*,x_*) - K(x_*,x)^T K_y^{-1} K(x_*,x) \tag{4}$$

$$K_y \triangleq K(x,x) + \sigma^2 I \tag{5}$$

The hyper-parameters, θ, can be optimised by maximising the log marginal likelihood given as (Murphy 2012):

$$\begin{aligned} \log p(y|x) &= \log \mathcal{N}\left(y|m,K\right) \\ &= -\frac{1}{2}\left((y-m)K^{-1}(y-m) + \log|K| + N\log(2\pi)\right) \end{aligned} \tag{6}$$

With the expression for the marginal likelihood (6) any standard gradient-based optimiser can be used to estimate the kernel parameters. Having defined the basic GP structure, we now proceed to elaborating our GP based model for water distribution system.

3 Noise Scaled GP for Leak Detection

Within a DMZ the consumption between consecutive weeks tends to follow a very similar periodic pattern (see Figure 1(a)). On a weekly scale the consumption at a particular time in a day can be modelled as a function of the historic readings at the same time in the previous week/s. Similarly the variations in the readings at a particular time in the week can also be modelled as a function of the historic variations (see Figure 1(b)). To incorporate this in the GP framework, we firstly define a weekly moving window on the data being modelled. Secondly, we define two vectors, the weekly mean (m_w) and weekly variance (σ_w^2). Assuming sensor data is discrete (with readings taken at fixed time intervals), the weekly means and variances vectors consist of a mean and variance value for each reading time slot, t, in a week i.e. $m_w(t)$ and $\sigma_w^2(t)$, where each mean and variance value is computed from the historic flow data using a moving average window of four weeks (see Figure 1(b)) as:

$$m_w(t) = \frac{1}{4}\sum_{i=1}^{4} y_i(t)$$

$$\sigma_w^2(t) = \frac{1}{4}\sum_{i=1}^{4}(y_i(t) - m_w(t))^2$$

The weekly mean vector can be easily modelled as a GP prior mean (Rasmussen and Williams 2005). However, heteroscedastic noise, in GP models results in a non-closed form likelihood expression. A tractable exception can be derived with the assumption of independence between noise variances (Venanzi, Rogers, and Jennings 2013). Given our model, we assume that the observed noise variance at time t is independent of the noise variance at t' i.e the values in the weekly variance vector are independent. This allows us to assume the noise at each time slot, t, to be Gaussian. The weekly variance, σ_w^2, is incorporated in our GP model as a diagonal noise matrix, $\Theta = \sigma_w^2 I$, which is added to the GP covariance structure $k(t,t;\theta_{cov})$ parameterised by θ_{cov}. Having defined the mean and covariance structure for the GP, the prior (1) and posterior predictive distribution (2) can be re-written as:

$$p(y|t,\theta_{cov}) = \mathcal{N}\left(y; m_w(t), k(t,t;\theta_{cov}) + \sigma_w^2(t)I\right) \tag{7}$$

$$p(y_*|t_*,t,y,\theta_{cov}) = \mathcal{N}\left(y_*; m(t_*|y,t), \Sigma(t_*|y,t,\theta_{cov})\right) \tag{8}$$

where

$$m(t_*|y,t) = m_w(t_*) + K(t_*,t)K_y^{-1}(y - m_w(t)) \quad (9)$$

$$\Sigma(t_*|y,t,\theta_{cov}) = K(t_*,t_*) - K(t_*,t)^T K_y^{-1} K(t_*,t) \quad (10)$$

$$K_y \triangleq K(t,t) + \Theta \quad (11)$$

$$\Theta = \sigma_w^2(t)I \quad (12)$$

Using diagonal noise matrix, Θ, inherently adds a bias to the diagonal of the covariance structure that reflects the actual variations in data and thus providing better handling of the input dependent noise in the underlying system.

4 Leak Detection and Quantification by Parameterizing the GP

We now return to the problem of leak detection and quantification. Leaks are non-deterministic and cause an increase in the measured flow thus changing the characteristic of the underlying system for the duration of the leak. This deviation from normal behaviour can be used to detect leaks in the system. Recent works on GP based fault/change-point detection (Garnett et al. 2010) propose various kernels to detect different types of faults. Out of these, of particular relevance to leak detection problem is the bias fault kernel (Garnett et al. 2010) which models faulty readings as a simple offset from the true values by some constant amount. We use the same approach to model a step change mean function, m_l, parameterized by $\theta_l = \{\theta_{loc}; \theta_{mag}\}$; where θ_{mag} models the magnitude of the leak and θ_{loc} represents the starting location of the leak, as:

$$m_l(t;\theta_l) = \left\{ \begin{array}{ll} \text{zero} & \text{if } t < \theta_{loc} \\ \theta_{mag} & \text{if } t \geq \theta_{loc} \end{array} \right\}$$

We incorporate this in our GP model by defining the GP prior mean as a composite mean function, m_c, which is a sum of the weekly mean and the step change mean function resulting in a semi-parametric GP (Murphy 2012). Given this equation (9) can be rewritten as:

$$m(t_*|y,t) = m_c(t_*) + K(t_*,t)K_y^{-1}(y - m_c(t)) \quad (13)$$

Since $m_c(t) = m_w(t) + m_l(t;\theta_l)$, we can train the GP on data without leaks by setting $m_l(t;\theta_l)$ to zero.

5 Hyperparameter Management

In this section we describe how the two sets of hyperparameters i.e. the covariance hyperparameters, θ_{cov}, and the mean hyperparameters, θ_l, are learned and how the two learning schemes are interwoven. In our model θ_{cov} represents the correlations in the DMZ flow/demand when there is no leak in the system. We use a squared exponential kernel, $kse(t,t') = \sigma_{out}^2 \exp(-(t-t')^2/2l^2)$, to model these correlations (for a discussion of kernel types see Rasmussen and Williams 2005). Thus in our model the hyperparameters in θ_{cov} are the length scale l and the output variance σ_{out}^2. We find the optimal values for these by setting $m_l(t;\theta_l)$ to zero and then using gradient decent search to train the GP on four weeks of aggregated flow data (Rasmussen and Williams 2005).

Once the GP is trained, for new observed flow values, detection is done over a moving window of one week by learning the mean function hyperparameters, $\theta_l = \{\theta_{loc}; \theta_{mag}\}$, using a bounded search process. At each time step, t, both θ_{mag} and θ_{loc} are bounded based on the observed flow. The leak location parameter, θ_{loc}, is bounded by our data modelling window of one week. For a new observed flow value, $y(t)$, at time, t, the leak magnitude will always be between $m_w(t)$ and $y(t)$. The leak at a particular time can not be more than the observed flow reading, $y(t)$, and less than the historic weekly mean, $m_w(t)$. An observed flow value less than the historic mean suggests a reduction in water flow where as a leak always produces an increase. Thus, at a particular time slot, t, if $y(t) - m_w(t) < 0$, then search for θ_{loc} at that particular time slot, t, can be pruned. These bounds reduce the leak parameter search space considerably. In case $y(t) - m_w(t) > 0$ or a set min threshold value, we discretize the possible θ_{mag} values by sampling five consecutive equidistance values between the upper and lower bounds, $y(t)$ and $m_w(t)$. Exhaustive search is then used to find both a single θ_{mag} (out of the discretized possibilities) and a θ_{loc} value (out of all possible time slots in a week) that yields the lowest negative log likelihood. It must be noted that by modelling the leak as a parametric mean function we have explicitly made the leak independent of the correlations in the observations and observation noise model under no leak conditions. This leads to a better modelling of the actual physical leak process as leaks are un-correlated additions to the underlying latent process. Furthermore, when learning the leak parameters, this results in the marginal likelihood giving an estimate of the leak based on a normal no leak observation noise model. Thus, any variations in the observation noise due to the leak are also captured in the leak parameters, giving a more accurate estimate of the leak effect. Having defined the leak detection process we now refer back to the previously mentioned requirement of distinguishing between legitimate short term increase in flow and leaks. Depending on the specific reliability requirements of a utility company, a duration, T, of continuous increased flow can be defined for an anomaly to be considered a leak. In such cases θ_{loc} will always be between the current timeslot t and $t - T$. This not only allows continuous monitoring and record keeping of the leak but also leak correction, in the observed flow data, based on any previously confirmed leaks.

6 Evaluation And Comparison on Real Water Network Data

To evaluate our model we use five weeks of real flow data from a DMZ in UK, with readings taken at 15 minutes interval. Four weeks of data is used to calculate the weekly mean and variance vectors. We introduce time varying leaks in the fifth week using the following formula.

$$L_s^n = B_{mag} + (B_{mag}/8)y_s^n$$

Where L represents the leak, s is the starting location of the leak, n is the total number of readings in a week, B_{mag} is a constant base magnitude of the leak and y is the observed flow. This allows us to simulate leaks that are dependant on the observed flow. In our experiments we consider the following metrics for evaluation:

1. **True Positive Rate (TPR):** This is the ratio of the number of correctly identified times slots with a leak, TP, over the true number of times slots with a leak P i.e. $TPR = TP/P$.
2. **False Positive Rate (FPR):** This is the ratio of the number of in-correctly identified times slots with a leak over, FP, the true number of times slots without a leak N i.e. $FPR = FP/N$.
3. **Accuracy (ACC):** Accuracy is the proportion of true results (both true positives and true negatives) in the population. It is is the degree of closeness of the number of detected leaks to the actual number of leaks in the data i.e. $ACC = (TP + TN)/(P + N)$.
4. **Precision (PRE):** Precision or positive predictive value is defined as the proportion of the true positives against all the positive results. It is measures the accuracy of the correct identifications of a leak. $PRE = (TP)/(TP + FP)$.
5. **Root Mean Square Error (RMSE):** To test the accuracy of the detected magnitude of a leak, we compute the root mean square error using the detected leak, $\hat{L_d}$, and the actual simulated leak, L_s as:

$$RMSE = \sqrt{\frac{1}{N}\Sigma_i^N(\hat{L_d} - L_s)^2}$$

We test the leak detection accuracy of our model in comparison to the recent GP based fault bucket method (Osborne et al. 2012) which is an extension of the previous work by Garnett et al.(2010). Furthermore, we also compare our model with two applied leak detection approaches in water industry. Firstly, leak detection by Nightline analysis (NL), which is the default choice leak detection approach widely used by water utility companies. Secondly, we compare our model with a Kalman filter based recent advancement in leak detection, which has been applied and tested in a real life scenario (Ye and Fenner 2010). We briefly describe each of the selected methods.

1. **Nightline Analysis (NL):** The the nightline flow for a day is defined as the minimum out of all the average hourly readings taken from 00:00 to 23:45 inclusive. The Nightline analysis algorithm models one nightline flow reading per day, i, as a discrete random variable $F_i = m_i + \epsilon_i$ where m_i is the mean and ϵ_i is assumed to be the normally distributed noise. To detect a leak hypothesis tests are conducted to see if the observed $\epsilon_i = F_i - m_{(i-1)}$ comes from the same distribution as $\epsilon_{(1 \text{ to i-1})}$. Since, the NL method only uses one minimum variance data point per day, the nightline flow, it can take up to two days to detect and report a leak.
2. **Neptune Project Kalman Filter (NKF):** As part of the Neptune project (Ye and Fenner 2010), a Kalman filter based approach was proposed to detect leaks. The NKF approach uses one Kalman filter per time slot in a week to model the weekly flow. Leaks are detected by analysing the normalised residuals at time t which are calculated by applying a moving averaging window on the data points over the last week.

$$\text{Residual(t)} = R(t) = \frac{F\{y(t)\} - F\{\hat{y}(t)\}}{F\{y(t)\}}$$

$$R(t) = \left\{ \begin{array}{ll} 0 & \text{if } R(t) < 0.01 \\ R(t) & \text{if } R(t) \geq 0.01 \end{array} \right\}$$

where F denotes the low pass filtering by the moving average window and 0.01 l/s is the minimum allowed leak. NKF can report leak as soon as they are detected. With observed flow readings 15 minutes apart the default leak reporting time for NKF is 15 minutes.
3. **Fault Bucket (FB):** The FB algorithm is based on the expectation that points that are more likely to be generated by noise with wide variance, than under the normal predictive model of a GP, are likely to be faults. This is formalised by choosing an observation noise distribution which is independent but not i.i.d as:

$$p(y|f, t, \neg\text{fault}, (\sigma^n)^2) = \mathcal{N}\left(y; f, (\sigma^n)^2\right)$$
$$p(y|f, t, \text{fault}, (\sigma^f)^2) = \mathcal{N}\left(y; f, (\sigma^f)^2\right)$$

where fault $\epsilon\{0, 1\}$ is an indication of fault presence or absence in observation $y(t)$ and $\sigma^f > \sigma^n$ is the standard deviation around the mean of the fault. Both σ^f and σ^n form part of the hyper-parameters of the model. The predictive distribution $p(f|y)$ for the latent variable f and the posterior probability of an observations faultiness $p(\sigma^f|y)$ are calculated by approximate marginalization based on four key assumptions, while the hyper-parameters are approximated by Bayesian Monte Carlo (for details see Osborne et al.(2012)). Since FB is an online detection method its default leak reporting time for our data set is also 15 minutes.

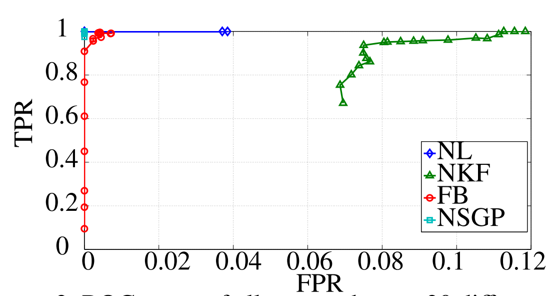

Figure 2: ROC curve of all approaches on 20 different leaks.

We test the detection rates of all algorithms over one week of observed data with a simulated time varying leak. This test is repeated with 20 leaks with different magnitudes (varying from 0.2 l/s to 1 l/s) starting at different locations within the week. For comparison with the GP based FB approach we use the standard squared exponential kernel in both FB and NSGP approaches with the same learned hyperparameters. The only difference being FB uses the standard constant Gaussian noise variance assumption where as in the NSGP we use the diagonal noise covariance kernel in conjunction with the proposed step change mean function. Figure 2 illustrates the performance of all algorithms using the Receiver Optimisation Characteristic Curve while

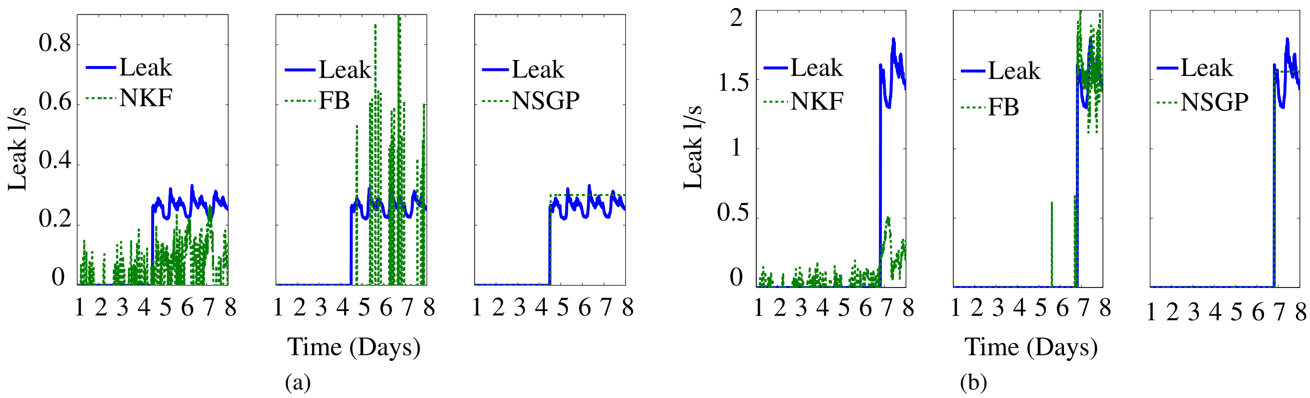

Figure 3: (a) Leak detection of all approaches for a leak starting at 12:00 PM on the 4th day with a base magnitude of 0.2 l/s. (b) Leak detection of all approaches for a leak starting at 7:30 PM on the 6th day with a base magnitude of 1.15 l/s

	Default Reporting Time	TPR	FPR	ACC	PRE	RMSE (l/s)	Computation Time (seconds)
		Averages over 20 different leaks					
NL	2 days	**1**	0.015	0.986	0.808	N/A	**0.011**
NKF	15 mins	0.915	0.221	0.813	0.650	0.397	0.068
FB	15 mins	0.811	0.001	0.911	0.995	0.153	60.204
NSGP	15 mins	0.997	**0.0**	**0.999**	**1**	**0.065**	63.276

Table 1: Detection results all approaches averaged over 20 different leaks

Table 1 details the results of the algorithms for each of the selected performance metric. Although it may seem that the NL algorithm performs better or at least as good as the proposed NSGP it must be noted that the NL algorithm addresses a much simpler problem, as it uses one minimum hourly average flow reading per day (nightline flow). Thus, even though the NL algorithm shows high performance on the nightline flow, it always takes at least 2 days for the NL algorithm to confirm and report a leak. Owing to this the NL algorithm is not included in Figure 3, which shows the leak detection results of the selected algorithms over two out of the 20 simulated leaks. However, the detection accuracy of the NL algorithm, using the nightline flow, has been compared with the remaining approaches. The results in Table 1 show that the proposed NSGP outperforms other approaches with a detection accuracy of 99%, the lowest false positive rate and the most accurate leak magnitude estimate with a RMSE of 0.0647 l/s. To elaborate the affects of Gaussian noise and leak independent assumptions (listed in Section 1), out of the 20 simulated leaks we have selected a small leak of 0.2 l/s, shown in Figure 3(a), which lies between the max, 0.4, and min, 0.0001, observed variance in data. Since leak detection in NKF is based on smoothed residuals and Gaussian noise assumptions, it does not cope well with the heteroscedasticity in the data, producing false positive detections before the actual leak occurrence (see Figure 3(a) and Figure 3(b)) resulting in the highest overall FPR rate of 0.221 and lowest ACC (0.813). The FB approach, like NKF, assumes a constant Gaussian noise variance and the leak at each new observed flow reading to be independent of the previous ones. As a consequence of these assumptions the FB approach produces false negatives for the small leak in Figure 3(a) even after the leak occurrence, resulting in the lowest TPR of 0.811 out of all the selected approaches. Owing to the false positive and false negative detections, both NKF and FB have higher leak magnitude estimation errors of 0.397 and 0.153 l/s respectively whereas NSGP gives the most accurate and sustained average estimate of the leak magnitude resulting in the lowest RMSE of 0.065 l/s. Although NSGP has the highest computation time, it is 18.6% more accurate than NKF and 8.8% more accurate than FB. It must be noted that, the computation cost of NSGP is dependent upon the location of the leak in the weekly window. In actual realtime deployment of NSGP leak correction would be applied to the data when leaks are confirmed after a set duration T, as discussed in Section 5. As a result, most of the locations from time slot $t = 1$ to (times slots in a week $-T$), would be pruned resulting in considerably lower computation times than in the current experiments without leak correction, where on average 68.45% of the locations were pruned.

7 Conclusion

We have proposed a novel GP based model to detect leaks in water distribution systems. Our chief contribution is domain specific application of GP to detect leaks in presence of heteroscedasticity. Also modelling of the physical leak process in a way that allows us to use domain knowledge to dynamically reduce the leak detection time while also catering for sustained leaks, reliability, and leak quantification requirements posed by the domain. As to future work we have deployed a prototype implementation of NSGP in the live systems at our research partner company I2O Water which we will evaluate in an online setting. We will also investigate methods to learn the correlations between multiple neighbouring DMZ to verify and distinguish between leaks and legitimate short term increase in consumption.

References

Chandola, V.; Banerjee, A.; and Kumar, V. 2009. Anomaly detection: A survey. *ACM Computing Surveys (CSUR)* 41(3):15.

Colombo, A.; Lee, P.; and Karney, B. 2009. A selective literature review of transient-based leak detection methods. *Journal of Hydro-environment Research* 2(4):212–227.

Consonni, G., and Marin, J.-M. 2007. Mean-field variational approximate bayesian inference for latent variable models. *Computational Statistics and Data Analysis* 52(2):790 – 798.

De Silva, D.; Mashford, J.; Burn, S.; Alliance., U. W. S. R.; for a Healthy Country Flagship (Program), W.; University., G.; of Queensland., U.; and CSIRO. 2011. *Computer aided leak location and sizing in pipe networks / Dhammika De Silva, John Mashford and Stewart Burn*. Urban Water Security Research Alliance, City East, Qld. :, 2nd ed. edition.

Garnett, R.; Osborne, M. A.; Reece, S.; Rogers, A.; and Roberts, S. J. 2010. Sequential bayesian prediction in the presence of changepoints and faults. *The Computer Journal* 53(9):1430–1446.

Goldberg, P. W.; Williams, C. K.; and Bishop, C. M. 1997. Regression with input-dependent noise: A gaussian process treatment. *Advances in neural information processing systems* 10:493–499.

Hwang, I.; Kim, S.; Kim, Y.; and Seah, C. E. 2010. A survey of fault detection, isolation, and reconfiguration methods. *Control Systems Technology, IEEE Transactions on* 18(3):636–653.

Kersting, K.; Plagemann, C.; Pfaff, P.; and Burgard, W. 2007. Most likely heteroscedastic gaussian process regression. In *Proceedings of the 24th international conference on Machine learning*, 393–400.

Kuss, M., and Rasmussen, C. E. 2005. Assessing approximations for gaussian process classification. In *Advances in Neural Information Processing Systems*, 699–706.

Lazaro-Gredilla, M., and Titsias, M. 2011. Variational heteroscedastic gaussian process regression. In Getoor, L., and Scheffer, T., eds., *Proceedings of the 28th International Conference on Machine Learning*, 841–848.

Murphy, K. 2012. *Machine learning a probabilistic perspective*. Cambridge, Mass: MIT Press.

Osborne, M.; Garnett, R.; Swersky, K.; and de Freitas, N. 2012. Prediction and fault detection of environmental signals with uncharacterised faults. In *26th AAAI Conference on Artificial Intelligence*, 533–540.

Osborne, M.; Garnett, R.; and Roberts, S. 2010. Active data selection for sensor networks with faults and changepoints. In *24th IEEE International Conference on Advanced Information Networking and Applications (AINA)*, 533–540.

Puust, R.; Kapelan, Z.; Savic, D.; and Koppel, T. 2010. A review of methods for leakage management in pipe networks. *Urban Water Journal* 7(1):25–45.

Rasmussen, C. E., and Williams, C. K. I. 2005. *Gaussian Processes for Machine Learning (Adaptive Computation and Machine Learning series)*. The MIT Press.

SWAN. 2011. Stated NRW (Non-Revenue Water) Rates in Urban Networks. Technical report, Smart Water Networks Forum (SWAN).

Venanzi, M.; Rogers, A.; and Jennings, N. R. 2013. Crowdsourcing spatial phenomena using trust-based heteroskedastic gaussian processes. In *First AAAI Conference on Human Computation and Crowdsourcing*.

Verde, C.; Morales-Menendez, R.; Garza, L.; Vargas, A.; Velasquez-Roug, P.; Rea, C.; Aparicio, C.; and De la Fuente, J. 2008. Multi-leak diagnosis in pipelines a comparison of approaches. In *MICAI'08. Seventh Mexican International Conference on Artificial Intelligence*, 352–357. IEEE.

Ye, G., and Fenner, R. A. 2010. Kalman filtering of hydraulic measurements for burst detection in water distribution systems. *Journal of Pipeline Systems Engineering and Practice* 2(1):14–22.

Zhang, P., and Ding, S. 2008. On fault detection in linear discrete-time, periodic, and sampled-data systems. *Journal of control science and engineering* 2008:8.