

Parametric Nonlinear Regression Models for Dike Monitoring Systems

Harm de Vries^{1,2,3}, George Azzopardi^{2,3}, André Koelewijn⁴, and Arno Knobbe¹

¹ Leiden Institute of Advanced Computer Science,
Leiden University, Leiden, The Netherlands

² TNO, Groningen, The Netherlands

³ Johann Bernoulli Institute for Mathematics and Computer Science,
University of Groningen, Groningen, The Netherlands

⁴ Deltares, Delft, The Netherlands

Abstract. Dike monitoring is crucial for protection against flooding disasters, an especially important topic in low countries, such as the Netherlands where many regions are below sea level. Recently, there has been growing interest in extending traditional dike monitoring by means of a sensor network. This paper presents a case study of a set of pore pressure sensors installed in a sea dike in Boston (UK), and which are continuously affected by water levels, the foremost influencing environmental factor. We estimate one-to-one relationships between a water height sensor and individual pore pressure sensors by parametric nonlinear regression models that are based on domain knowledge. We demonstrate the effectiveness of the proposed method by the high goodness of fits we obtain on real test data. Furthermore, we show how the proposed models can be used for the detection of anomalies.

Keywords: Structural health monitoring, dike monitoring, nonlinear regression, anomaly detection.

1 Introduction

Dikes are artificial walls that protect an often densely populated hinterland against flooding disasters. Especially the Netherlands, with large areas below sea level, has a rich history of dike failures that resulted in drowning deaths and devastation of infrastructure. Although dike technology has improved over the years, only 44% of the 2875 kilometers of main Dutch dikes met the government's dike regulations in 2006 [1]. Traditional dike monitoring involves visual inspection by a dike expert at regular time intervals. Dike patrolling is, however, a time consuming and costly process that does not always reveal weak spots of a dike. In this light, the IJkdijk foundation¹ has been established in 2007 with the ambition to enhance dike monitoring by sensor systems. The largely successful program initiated an EU-funded project, called UrbanFlood², that also implements sensor systems in dikes, but intends to construct so-called Early Warning

¹ Official website: <http://www.ijkdijk.nl>

² Official website: <http://www.urbanflood.eu>

Systems (EWS) for floodings [2]. As a consequence of both projects, more and more dikes across Europe are being equipped with sensor systems, and therefore there is an urgent need for algorithms that are capable of detecting damage to the dike as early as possible.

Geophysical models [3,4] can be used to assess the stability of a dike by simulation of the underlying physical processes. Such models are computationally intensive and thus not appropriate for (near) real-time dike monitoring. To overcome such problems, data driven techniques were considered to detect indicators for instability of a dike. In [2], the authors proposed neural clouds in order to detect outliers in the sensor values. The main drawback of this approach is that sensor values are highly influenced by environmental conditions, and thus outliers often correspond to rare environmental conditions rather than changes in the internal structure of the dike. The same research group also proposed an anomaly detection technique that uses one-step-ahead prediction of (non-linear) autoregressive models [5]. Although such methods achieve high model fits, they are not appropriate to detect gradual changes in a response of a sensor.

In this paper, we conjecture that in order to detect internal changes in a dike, we first have to model the *normal* relationship between some environmental conditions and dike sensors. We present a case study of a set of pore pressure sensors that are installed in a sea dike in Boston (UK), and which are continuously affected by water levels, the foremost influencing environmental factor. Although the pore pressure signals vary significantly among the set of sensors, we hypothesize that essentially two physical processes play a role. We estimate the one-to-one relationships by parametric nonlinear regression models that aim to reflect the underlying physical phenomena. In contrast to black box modeling techniques, such as Transfer Function Modeling and Neural Networks, the proposed models are intuitive, interpretable and provide more insight into the dynamics of the dike. Moreover, we demonstrate that the models can be effectively used for anomaly detection.

2 Background

Fig. 1 shows a schematic overview of the concerned dike in Boston (UK) that includes the placement of seven sensors that measure the pore pressure at time intervals of 15 minutes. Although not shown in Fig. 1, there is a sensor that measures the water level nearby the dike with a sample interval of 15 minutes. In Fig. 2b, we show the water levels of the month of October 2011. It is characteristic of the dike in Boston that water levels follow half-daily tides with extreme differences (up to seven meters) between high and low tides. Note that the amplitude of the half-daily tides also varies with an approximately two-weekly period due to the lunar cycle. It is also worth mentioning that the sensor cannot detect water levels below -1.6 meters, which is reflected in the data by the flat lower envelope of the water level signal. The estuary near the dike falls dry at that point, although the actual sea levels are a little bit lower.

In Fig. 2a, we illustrate all seven pore pressure signals recorded in October 2011. The relationship between a pressure signal and a water level signal is

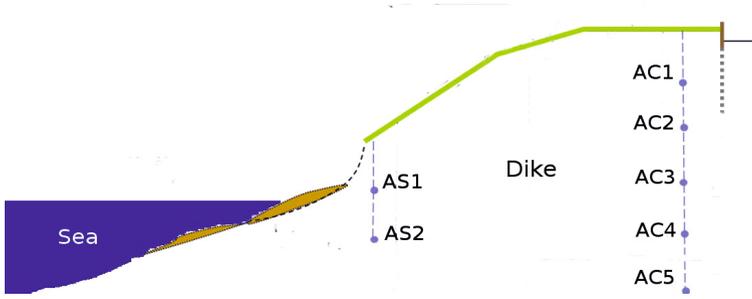


Fig. 1. Schematic overview of the sensor setup in the Boston (UK) dike

influenced by the location of the concerned pressure sensor. For instance, the AC1 sensor, which is placed at the top of the dike, does not respond to the water level. On the other hand, the AC4 pore pressure, which is located at a very deep level, seems to follow the same tidal fluctuations as the water levels.

The available data set consists of one year of sensor values, and ranges from October 2011 till October 2012. The data set has a lot of missing values. In particular, approximately 10% of the water levels and 20% of the pore pressure values are missing. Moreover, measurements of different sensors are not synchronized. As a preprocessing step, we therefore linearly interpolate the water level signal in order to align it with the pore pressure signal in question. In this way, we also fill small gaps in the water level signal of at most 2 samples (i.e. 30 minutes). The models we propose use some history of water levels to model the current pore pressure. We exclude the sample from the training set if either the pore pressure value is missing, or there is a gap of at least two water level measurements in the history. In Fig. 2, a sequence of missing values is visible as a straight line.

3 Model Estimation

In general, we expect that two physical phenomena play a role in the response of a pore pressure sensor:

Short-term effect. An almost immediate response to the water levels due to water pressure at the front of the dike. Therefore, the regular rise and fall of the water levels cause peaks in the pore pressure signals.

Long-term effect. A much slower effect that accounts for the degree of saturation of the dike. A dike that is exposed to high water levels absorbs water, which increases the degree of saturation of the dike, which in turn increases the pore pressure.

Fig. 2b clearly shows that the long-term effect does not play a role in all sensors. For example, the sensors AC4 and AC5, that reside deep in the dike, do not

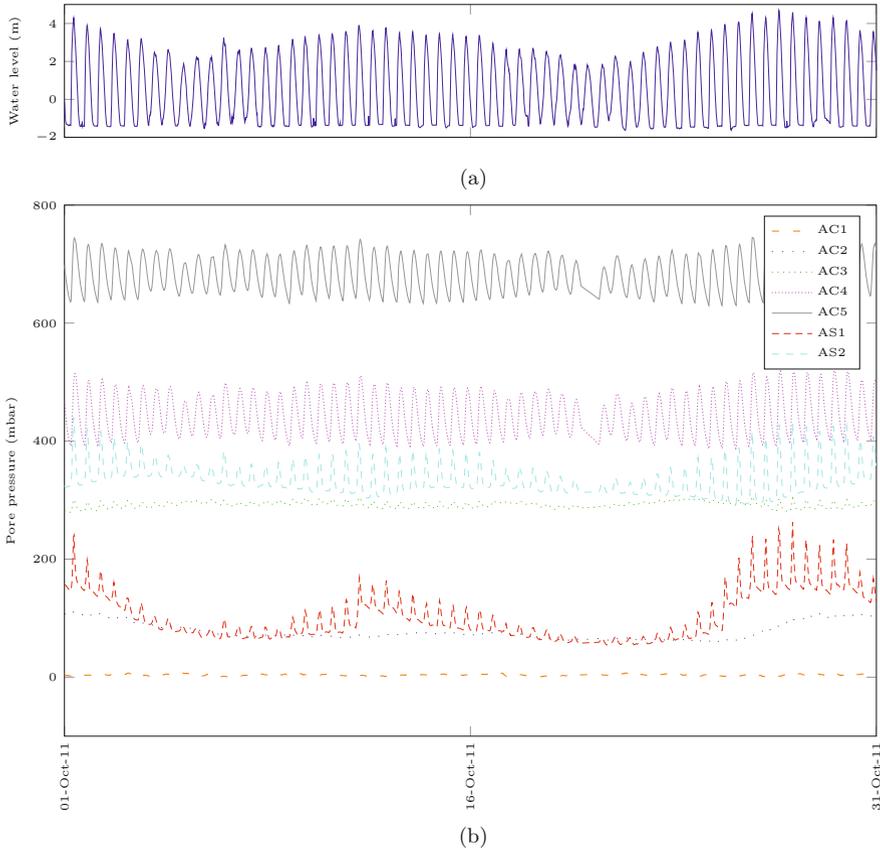


Fig. 2. The raw signals of one month of data of (a) the water level sensor and (b) the seven pore pressure sensors installed in the dike in Boston (UK). The straight lines indicate missing values.

have significant changes in the baseline. The saturation degree at that location in the dike is not heavily influenced by the dike's exposure to water levels. In the following, we refer to these sensors as *short-term effect sensors*, and we propose a model for them in Section 3.1.

The other sensors follow a mixture of both effects. A typical example is the AS1 sensor, of which an example month of data is shown in Fig. 2b. There are significant changes in the baseline of the signal that seem to follow the two-weekly cycle of the water levels, but there are still half-daily peaks that are superimposed on this baseline. We model these *mixed effect sensors* in Section 3.2.

3.1 Short-Term Effect Sensors

In the following, we select a subset of the complete data set as training data. Formally, we consider N pore pressure values $y[1], \dots, y[N]$ (equally spaced in time) and water level measurements $x[1], \dots, x[N]$ that are aligned in time. We model the current pore pressure value $y[n]$ as a function of recent history of water levels $x[n - M + 1], \dots, x[n]$. Here, $M > 0$ represents the number of historical water level measurements. It should be chosen large enough in order to reliably predict the pore pressure. We choose $M = 100$, which translates to roughly one day of water levels. The training set contains $P = N - M + 1$ examples, and consists of a set of water level input vectors $X = \{\mathbf{x}_i = x[i + M - 1], \dots, x[i] \mid i = 1, \dots, P\}$ and pore pressure output $\mathbf{y} = \{y_i = y[i + M - 1] \mid i = 1, \dots, P\}$.

We propose to model the short-term effect sensor by the following parametric nonlinear regression model³:

$$f(\boldsymbol{\theta}; \mathbf{x}_i) = b + a \sum_{m=0}^{M-1} \exp(-\lambda m) \mathbf{x}_i[m + 1] \quad \text{with } \boldsymbol{\theta} = [b \ a \ \lambda] \quad (1)$$

where λ controls the rate of decay, and a and b are affine transformation parameters. Our model corresponds to the solution of a first-order constant coefficient differential equation⁴:

$$y'(t) = -\lambda y(t) + ax(t) \quad \text{with initial condition } y(0) = b \quad (2)$$

By rewriting the right hand side to $\lambda(cx(t) - y(t))$ with $c = \frac{a}{\lambda}$, our model assumptions become clear. First, the water level linearly relates (by factor c) to the pressure on the front of the dike. Second, the rate of change of the pore pressure is proportional to the difference between the current pressure on the front of the dike and the pore pressure of the sensor.

We estimate the parameters of the model by minimizing the sum of squared residuals:

$$S(\boldsymbol{\theta}) = \sum_{i=1}^P r_i^2(\boldsymbol{\theta}) \quad \text{where } r_i(\boldsymbol{\theta}) = y_i - f(\mathbf{x}_i; \boldsymbol{\theta}). \quad (3)$$

which is identical to the Maximum Likelihood Estimator (MLE) under white Gaussian error terms. We optimize the cost function by a Gauss-Newton solver, which is appropriate to optimize a least-squares problem [6].

We fit the model on 12 days of AC4 sensor values. We obtain the parameters $b_{mle} = 430.13$, $a_{mle} = 3.17$ and $\lambda_{mle} = 0.1142$, and show the predicted values by

³ For initial rest (i.e. $b = 0$) the proposed model is a Linear Time Invariant (LTI) system with an exponential decaying impulse response function. Linear constant coefficient differential equations can be represented by causal LTI systems if and only if they satisfy the initial rest condition. We refer the interested reader to http://ocw.mit.edu/resources/res-6-007-signals-and-systems-spring-2011/lecture-notes/MITRES_6_007S11_lec06.pdf.

⁴ See <http://web.mit.edu/alecp/www/useful/18.03/Supplementary-CG.pdf> for more details.

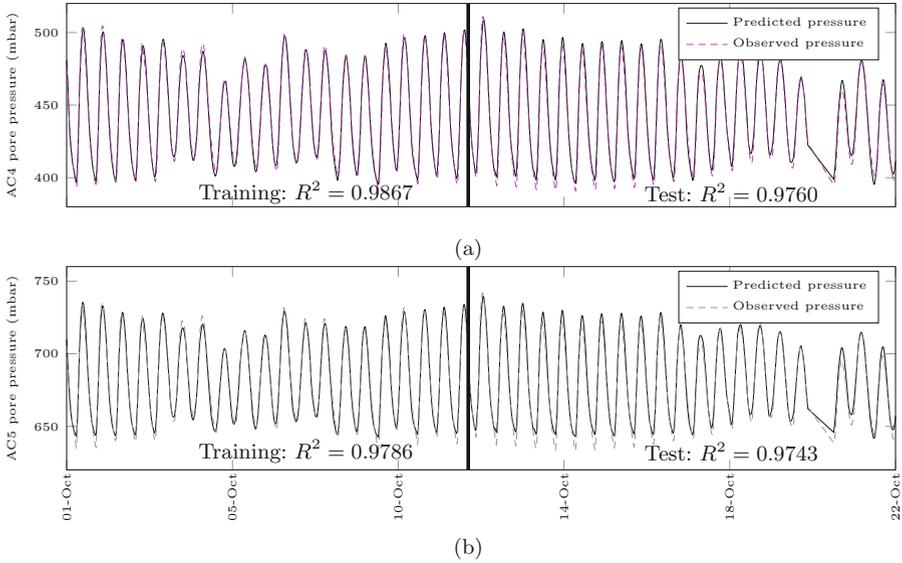


Fig. 3. A comparison between the predicted and observed values for (a) the AC4 and (b) the AC5 sensors. The black bar separates the training and test sets.

our model in Fig. 3a. Note that the model almost perfectly follows the observed pore pressure. We quantify the goodness of fit by:

$$R^2 = 1 - \frac{\sum_{i=1}^P (y_i - f(\theta; \mathbf{x}_i))^2}{\sum_{i=1}^P (y_i - \bar{y})^2} \quad \text{with mean } \bar{y} = \sum_{i=1}^P \frac{y_i}{P}, \quad (4)$$

which, roughly speaking, measures how successful the model is in explaining the variation of the data. For the above example, we find $R^2 = 0.9867$, which indicates that the estimated model fits very well. We consider as test data the 12 days that follow the training data. For this period we also obtain a very high value of $R^2 = 0.9760$, which demonstrates that the proposed regression model is not prone to over-fitting.

To give an impression of the duration of the exponential decay, note that $\lambda = 0.1142$ corresponds to a mean lifetime of $\tau = 1/0.1142 = 8.756$ measurements, which in our case amounts to slightly over 2 hours. The corresponding half-life is $\tau_{1/2} = 8.756/\ln(2) = 6.070$ which corresponds to roughly 1.5 hours.

In a similar way, we estimate the parameters for the model on 12 days of AC5 pore pressure values, and obtain $b_{mle} = 671.78$, $a_{mle} = 3.09$ and $\lambda_{mle} = 0.1458$. We show the predicted values obtained by our model in Fig. 3b, and point out that the predicted values are quite close to the observed pore pressure measurements. This is confirmed by the goodness of fit, $R^2 = 0.9786$. For the next 12 days, we obtain $R^2 = 0.9743$, which demonstrates its effectiveness on unseen data.

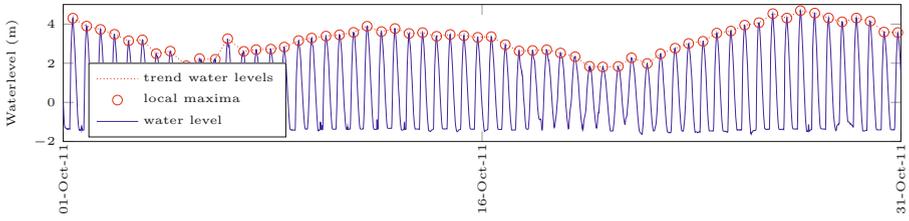


Fig. 4. Extraction of the top envelope of the water levels. The circles indicate the local maxima which are connected by straight lines using linear interpolation.

3.2 Mixed-Effect Sensors

We extend the models that we proposed above to sensors that are also affected by the degree of saturation of the dike. Intuitively speaking, a dike is only saturated if it is exposed to high water levels for a longer period. That is to say, we assume that the process varies slowly and it is not influenced by half-daily tides, but it is related to the two-weekly cycle of the water levels.

We extract the top envelope of the water levels in order to capture the general trend (the underlying fortnightly cycle). To this end, we extract the local maxima by taking into account the fact that local maxima are expected to be separated by 12 hours. Fig. 4 shows the extracted local maxima for the water levels in October 2011. This extraction decreases the resolution of the signal considerably. We use linear interpolation to fill in the gaps between the extracted local maxima, such that we obtain the same resolution as the given water level signal. We denote by $\hat{x}[1], \dots, \hat{x}[N]$ the estimated general trend of the water levels, and plot it in Fig. 4.

We model the long-term effect with a first-order constant coefficient differential equation that is similar to the short-term model defined in Eq. 2, but here we use the general trend of the water levels $\hat{\mathbf{x}}$ as input signal. The underlying assumption is that the rate of change in saturation degree is proportional to the difference between the current trend in water level (i.e. how much water the dike is currently exposed to) and the current saturation degree. In other words, a dike that is exposed to high water levels absorbs water much faster whenever it is not saturated. We propose to model the mixed-effect sensors by superimposing the short-term model on the long-term one:

$$\begin{aligned}
 f(\boldsymbol{\theta}; \mathbf{x}, \hat{\mathbf{x}}) = & b + a_l \sum_{m_l=0}^{M_l-1} \exp(-\lambda_l m_l) \hat{\mathbf{x}}[m_l + 1]^{p_l} \\
 & + a_s \sum_{m_s=0}^{M_s-1} \exp(-\lambda_s m_s) \mathbf{x}[m_s + 1]^{p_s}, \quad (5)
 \end{aligned}$$

where $\boldsymbol{\theta} = [b \ a_l \ \lambda_l \ p_l \ a_s \ \lambda_s \ p_s]^\top$, and subscripts s and l indicates the variables for short- and long-term effect, respectively. Note that we have included new

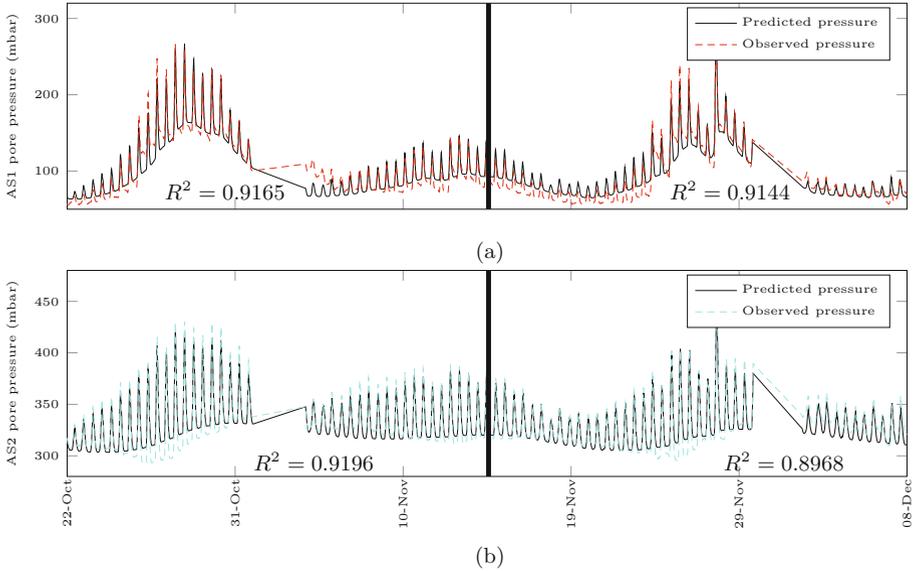


Fig. 5. A comparison between predicted and observed pore pressure values for (a) the AS1 sensor and (b) the AS2 sensor. The black bar separates the training and test sets.

parameters p_l and p_s that raise the water levels $\hat{\mathbf{x}}$ and \mathbf{x} to the powers of p_l and p_s , respectively. This means that the relation between water level (general trend of water level) and the pressure (saturation degree) on the front of the dike is assumed to be $c(x[n])^p$. Exploration of the sensor signals revealed that the response of some sensors is much higher to water levels above a particular threshold. There are several explanations for such an effect, which include the vertical height of the sensor that is below or above a particular water level, a change in slope of the dike front or a change in material covering the dike around that height. From the actual setup of the dike, as shown in Fig. 1, it is not immediately clear which of these the underlying reason might be.

In general, we expect that the long-term effect is based on a longer history than the short-term effect; i.e. $M_l \gg M_s$. Here, we set $M_l = 2000$ and $M_s = 100$, which correspond to roughly 21 and 1 day(s), respectively. The mixed-effect model is more richly parameterized than the short-effect model, and therefore we also need more data to reliably estimate the parameters. For the mixed-effect sensor, the training set is of size $P = N - M_l + 1$, and we decided to use approximately 21 days of training data; i.e. $P = 2000$. The training set consists of the general trend of water level input vectors $\hat{X} = \{\hat{\mathbf{x}}_i = x[i + M_l - 1], \dots, x[i] \mid i = 1, \dots, P\}$, water levels input vectors $X = \{\mathbf{x}_i = x[i + M_l - 1], \dots, x[i + M_l - M_s] \mid i = 1, \dots, P\}$, and pore pressure output $\mathbf{y} = \{y_i = y[i + M_l - 1] \mid i = 1, \dots, P\}$.

Table 1. The estimated parameters for the mixed-term sensors, as well as the goodness of fit R^2 for training and unseen data

Sensor	b	a_l	λ_l	p_l	a_s	λ_s	p_s	R^2	R_{unseen}^2
AS1	57.75	$4.32e-11$	0.0089	10.59	$2.92e-6$	0.83	7.47	0.9165	0.9144
AS2	283.31	$6.18e-11$	$7.15e-4$	9.41	$2.2e-3$	1.24	4.60	0.9196	0.8968

We estimate the set of parameters θ by MLE, which minimizes the sum of squared residuals as defined in Eq. 3, but this time the residuals $r_i(\theta) = y_i - f(\theta; \mathbf{x}_i, \hat{\mathbf{x}}_i)$ are defined with respect to the new model $f(\theta; \mathbf{x}, \hat{\mathbf{x}})$. Table 1 presents the estimated parameters for the AS1 and AS2 sensors, as well as the goodness of fit for training and test data. The values determined by the proposed models for training and test sets are shown in Fig. 5. The AS1 sensor model only partly captures the long-term effect. In particular, there is a significant difference around November 3-4. The short-term effect is also not modeled very accurately, but this might be a consequence of the imperfect long-term effect model. Our observation is confirmed by the goodness of fit $R^2 = 0.9165$, which is slightly worse than the goodness of fit for short-term effect sensors. For unseen data we obtain $R^2 = 0.9144$, which indicates that the model captures at least some of the underlying dynamics of the AS1 sensor.

The AS2 sensor is dominated by the short-term effect, and only has a minor contribution from the long-term one. Fig. 5b shows that the estimated model captures the short-term effect, but fails to learn the long-term one. Nevertheless, the goodness of fit for both training and test set is in the order of $R^2 = 0.91$.

We have excluded the results of the AC2 and AC3 sensors, since they are in line with the AS2 sensor; the proposed model is not robust enough to fully capture the long-term effect. We believe that this is mainly due to other environmental factors (e.g. outside temperature, humidity, and air pressure) that were not considered (because they were not available) in the proposed model.

4 Anomaly Detection

In the previous section, we showed that we can reliably estimate a model for the short-term effect sensors. As an example application, we show that such models can be effectively used to detect changes in the response of the sensor — the so-called anomalies.

We employ the following semi-supervised strategy to detect anomalies in the AC4 sensor. We first estimate the parameter set θ of the model on data of October 2011 (that is explicitly labeled as normal). We then use the estimated model $f(\theta; \mathbf{x})$ to predict future pore pressure values $\tilde{y}[n]$ for the next months. A pointwise anomaly score is calculated by measuring the Euclidean distance $\sqrt{(y[n] - \tilde{y}[n])^2}$ between the predicted and observed pore pressure. We plot the anomaly score of the AC4 sensor till August 2012 in Fig. 6a, and mark two anomalies in this plot by colored rectangles. The red rectangle indicates a rather

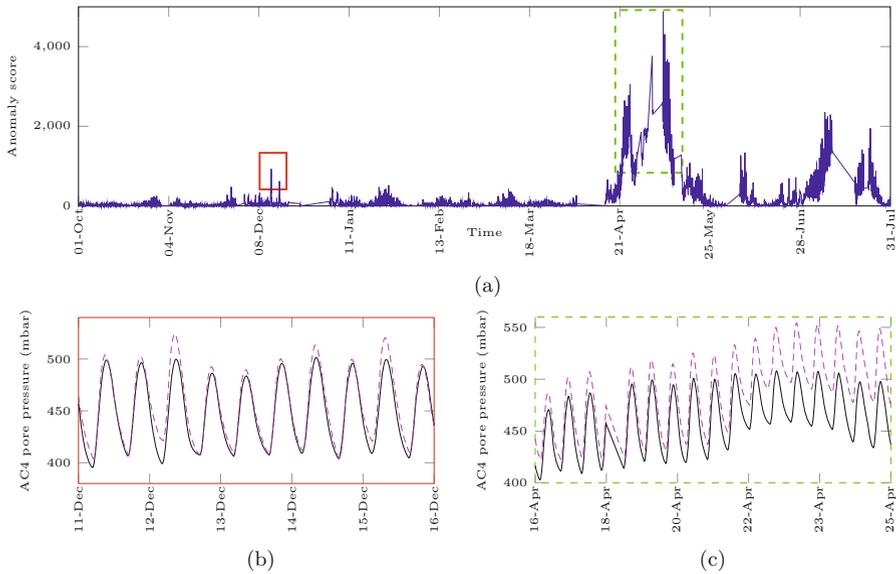


Fig. 6. (a) Anomaly score for the AC4 sensor that is computed as the pointwise Euclidean distance between the predicted and observed pore pressure measurements. (b-c) A zoomed-in visualization of the predicted (purple dashed) and observed (black solid) pressure measurements for the respective red solid and green dashed marked regions in (a).

small anomaly that corresponds to a small offset between the predicted and observed pressure around December 12, Fig. 6b. On the other hand, the green rectangle indicates a more serious anomaly since the anomaly score is high over a long period of time. Indeed, Fig. 6c illustrates a significant discrepancy between predicted and observed pore pressure from April 16 till April 25. Note that the anomaly score returns to almost zero around half of May, and thus the detected anomaly is not a structural change in the response of the sensor. April 2012 was characterized by extreme rainfall⁵, and we speculate that the anomaly is caused by outflow from the locks just upstream of the monitoring site.

5 Discussion and Conclusions

In this paper, we proposed parametric nonlinear regression models that describe the relationship between a water height sensor and individual pore pressure sensors. The models that we propose are highly effective (in the order of goodness of fit $R^2 = 0.97$) for pore pressure sensors (AC4 and AC5) that exhibit short-term physical phenomenon. Moreover, we demonstrated that the proposed model can be effectively used for the detection of anomalies.

⁵ See <http://www.metoffice.gov.uk/climate/uk/summaries/2012/april>

While we also achieved reasonably high goodness of fit for what we refer to as long-term effect sensors (AS1 and AS2), we believe that the proposed model can be further enriched by incorporating information about other environmental factors, such as rainfall, humidity and outside temperature. Although the presented models are tailored to sensors installed in a sea dike, we think that, due to the general nature of the applied techniques, they are applicable to other sensing and monitoring systems.

Acknowledgement. This work has been funded by TNO. We thank Jeroen Broekhuijsen for making the data available and for valuable discussions.

References

1. Van Baars, S., Van Kempen, I.: The causes and mechanisms of historical dike failures in the Netherlands (2009)
2. Pyayt, A., Mokhov, I., Lang, B., Krzhizhanovskaya, V., Meijer, R.: Machine learning methods for environmental monitoring and flood protection. *World Academy of Science, Engineering and Technology* (78), 118–124 (2011)
3. Melnikova, N., Shirshov, G., Krzhizhanovskaya, V.: Virtual dike: multiscale simulation of dike stability. *Procedia Computer Science* 4(0), 791–800 (2011); *Proceedings of the International Conference on Computational Science* (2011)
4. Tua, P., van Gelder, P., Vrijlinga, J., Thub, T.: Reliability-based analysis of river dikes during flood waves
5. Pyayt, A., Kozionov, A., Kuserbaeva, V., Mokhov, I., Krzhizhanovskaya, V., Broekhuijsen, B., Meijer, R., Slood, P.: Signal analysis and anomaly detection for flood early warning systems (2014)
6. Nocedal, J., Wright, S.: *Numerical optimization*, vol. 2. Springer, New York (1999)