

Journal Pre-proofs

Research papers

Shallow Aquifer Monitoring Using Handpump Vibration Data

Achut Manandhar, Heloise Greeff, Patrick Thomson, Rob Hope, David A Clifton

PII: S2589-9155(20)30008-0
DOI: <https://doi.org/10.1016/j.hydroa.2020.100057>
Reference: HYDROA 100057

To appear in: *Journal of Hydrology X*


Received Date: 22 October 2019

Accepted Date: 18 May 2020

Please cite this article as: A. Manandhar, H. Greeff, P. Thomson, R. Hope, D.A. Clifton, Shallow Aquifer Monitoring Using Handpump Vibration Data, *Journal of Hydrology X* (2020), doi: <https://doi.org/10.1016/j.hydroa.2020.100057>

This is a PDF file of an article that has undergone enhancements after acceptance, such as the addition of a cover page and metadata, and formatting for readability, but it is not yet the definitive version of record. This version will undergo additional copyediting, typesetting and review before it is published in its final form, but we are providing this version to give early visibility of the article. Please note that, during the production process, errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

© 2020 Published by Elsevier B.V.



Shallow Aquifer Monitoring Using Handpump Vibration Data

Achut Manandhar^a, Heloise Greeff^a, Patrick Thomson^b, Rob Hope^b, David A Clifton^a

^a*Department of Engineering Science, University of Oxford, UK*

^b*School of Geography and the Environment & Smith School of Enterprise and the Environment, University of Oxford, UK*

Abstract

We present a novel technology for monitoring changes in aquifer depth using handpump vibration data. This builds on our previous works using data to track handpump usage and facilitate handpump maintenance systems in rural parts of Kenya. Our motivation is to develop a cost-effective and scalable infrastructure to monitor shallow aquifers in regions where handpumps are already part of water infrastructure, but where traditional sources of groundwater monitoring data may be limited or non-existent. The data is generated using accelerometer sensors attached to the handles of nine handpumps in the study site in Kenya, instrumented for a year. These time-series data from handpumps are individually modelled using machine learning methods to track the changes in the water level with respect to the bottom of the rising main. Results show promise in modelling handpump vibration data with machine learning approaches to provide useful aquifer monitoring information from the “accidental infrastructure” of community handpumps. This technology is intended to complement existing hydrogeological modelling, and one of our key future goals is to integrate these machine learning out-

puts with hydrogeological information to develop more refined and robust models for shallow aquifer monitoring.

Keywords:

Aquifer monitoring, Water column, Handpumps, Machine learning, LSTM, Novel techniques

1. Introduction

Groundwater is directly linked to United Nations' Sustainable Development Goal 6 (SDG 6) - clean water and sanitation for all by 2030 (UN, 2018). It is also inexplicably linked to other SDGs related to poverty eradication, food security, gender equality, sustainability of cities and human settlement, combating climate change, and protecting terrestrial ecosystems (Hope and Rouse, 2013; Howard, 2014; UN, 2018; WWAP, 2019). It is estimated that groundwater provides around 50% of all drinking water and 40% of all agricultural irrigation worldwide (FAO, 2011). In Africa, groundwater is the major source of drinking water and its use for irrigation is expected to increase substantially to tackle growing food insecurity (MacDonald et al., 2012).

The magnitude of groundwater's significance is in sharp contrast to the dearth of reliable quantitative information on groundwater resources, especially in Africa (MacDonald et al., 2012; Fan et al., 2013; UN, 2018). Available data on global groundwater storage estimates, often based on decade-old studies with large uncertainties, limit effective governance of groundwater systems (Richey et al., 2015). Like any other resources, groundwater resource management requires monitoring to make informed plans by water users and

20 water service providers at regional and national levels. However, long-term
 21 monitoring data are often scarce in Africa, and wherever data are available,
 22 inconsistencies in methodologies used at different times make comparisons
 23 difficult and trend unclear (Comte et al., 2016). Traditional groundwater
 24 monitoring technologies (Xu et al., 2012; Van Camp et al., 2013; Kelbe et al.,
 25 2016) are often data, resource, and time intensive. Recent efforts have shown
 26 remote sensing observations can provide useful cost-effective auxiliary data
 27 to improve global groundwater estimates (Richey et al., 2015). We propose a
 28 novel complementary shallow aquifer monitoring technology that utilizes the
 29 continent’s existing handpump infrastructure. Handpumps remain a reliable
 30 and low-cost method to access groundwater in the context of rural water
 31 supply for around 200 million people in Sub-Saharan Africa (Thomson et al.,
 32 2012). We aim to explore if a network of these handpumps can potentially
 33 provide auxiliary information that can be exploited using machine learning
 34 approaches to monitor the underlying shallow aquifer systems.

35 Previous efforts have shown that vibration data collected at the hand-
 36 pump’s handle are indicative of pump malfunction (Thomson et al., 2012;
 37 Greeff et al., 2019). Changes in the characteristics of vibration data, poten-
 38 tially due to handpump malfunction, can be tracked using novelty detection
 39 approaches. Higher novelty scores signify departure from normal operating
 40 conditions. Remote transmission of these novelty scores, as part of a hand-
 41 pump maintenance infrastructure, can be used for rapid pump maintenance.

42 The changes in vibration data were shown to be indicative of changes in
 43 the water level at the borehole under controlled circumstances (Colchester
 44 et al., 2014, 2017). These works showed that the vibration generated at the

45 handpump's handle are affected by the weight of the system, i.e. the weight
 46 of the rods and the volume of water inside the rising main. In general,
 47 the deeper the water level, the higher the volume of water inside the rising
 48 main, and the stronger the observed vibration. The variation in vibration
 49 characteristics was exploited to estimate the water level at the borehole of
 50 the handpump. These experiments were performed using two datasets -
 51 (1) data from a handpump on the University site, primarily installed to
 52 perform condition monitoring related research, and (2) few weeks of data
 53 from handpumps in the study site in Kenya. Both datasets were collected
 54 under controlled settings, i.e. handpump users were mostly limited to the
 55 research team, and the pump strokes and duration were fairly consistent.
 56 Given the short data collection period, the dataset neither captured long-
 57 term variation in the aquifer level nor faced challenges commonly encountered
 58 during normal operations, e.g. wide variation in the pumping style, variation
 59 in the water column at different times of the day, handpump breakdown
 60 and subsequent repairs, etc. Besides, for the dataset from the Kenya study
 61 site, instead of actually measuring the varying water column, the handpump
 62 rod length was used as an approximate measure of the depth of water in
 63 the borehole. The preliminary results of these experiments motivated this
 64 follow-up field study to determine if (1) the vibration data obtained from
 65 community handpumps in an unconstrained real-world setting can be used
 66 to track long-term changes in aquifer level (a conditional yes), and (2) the
 67 results generalize across different depths of shallow aquifer systems (yes).

68 Provided a handpump maintenance infrastructure is already in place,
 69 there is an opportunity to dual purpose the vibration data to monitor shallow

70 aquifer at minimal extra cost. The distributed maintenance infrastructure
 71 (Greeff et al., 2019) allows adaptive transmission of handpump vibration
 72 data, controlling when and how much of data is transmitted, thus saving
 73 cost in resource-constrained settings. The adaptive framework blends in well
 74 with the proposed aquifer monitoring system, whose predictions are only
 75 expected to be based on data from normally operational handpumps. This
 76 framework is further detailed in Section 2.

77 *1.1. Related works and our contribution*

78 Machine learning approaches have been explored to understand hydro-
 79 logical processes and facilitate water resource management (Herrera et al.,
 80 2010; Yaseen et al., 2019; Achieng, 2019). These approaches have also been
 81 used to model variations in aquifer level based on hydro-climatic data (e.g.,
 82 rainfall, temperature) (Nayak et al., 2006; Behzad et al., 2010; Yoon et al.,
 83 2011; Taormina et al., 2012; Tapoglou et al., 2014; Suryanarayana et al.,
 84 2014; Sahoo et al., 2017). These efforts have mostly implemented stan-
 85 dard regression approaches, e.g. Support Vector Regression (SVR) (Bishop,
 86 2006), Gaussian Processes (GPs) (Rasmussen, 2004), Multi-layer Perceptron
 87 Regression (MLP Regression) (Robert, 2014), or various hybrid combina-
 88 tions of the same. A recent work has shown Long Short Term Memory
 89 networks (LSTMs) outperformed standard MLP regression to predict water
 90 table depth by using 14 years of monthly water diversion, evaporation, pre-
 91 cipitation, temperature, and time as input data (Zhang et al., 2018). These
 92 efforts show the potential of using machine learning approaches to predict
 93 long-term changes in aquifer level in areas where hydrogeological data are
 94 difficult or expensive to obtain.

95 The proposed work also compares LSTMs with MLP regression for aquifer
 96 prediction. However, unlike all prior work, the proposed framework is novel
 97 because (1) it uses handpump vibration data to model changes in water level,
 98 and (2) setting aside the specific types of models used, the proposed frame-
 99 work combines regression and novelty detection approaches to develop a novel
 100 shallow aquifer monitoring technology that is designed to work alongside a
 101 handpump maintenance infrastructure. The proposed framework can also be
 102 extended to incorporate hydro-climatic data or outputs from hydrogeological
 103 models.

104 In an effort to facilitate reproducibility and data reuse, the dataset (Greeff
 105 et al., 2020) ¹ and the software (Manandhar, 2019) are both made publicly
 106 available.

107 2. Methods

108 In this work, we use a machine learning approach to model water column
 109 variation in a borehole based on vibration data observed at the handle of the
 110 handpump. We note that there are two types of changes in the water level
 111 – (i) the daily change in water level due to daily drawdown and recovery
 112 and (ii) the long-term seasonal change in water level due to aquifer discharge
 113 and recharge. In this work, we are interested in the latter. Although mod-
 114 elling the daily cycle is in itself an interesting problem, the challenge here is
 115 the unavailability of sufficient relevant vibration data because the drawdown
 116 occurs rapidly (and is often associated with pump priming whose vibration

¹The dataset will be available at the cited address once the embargo is lifted.

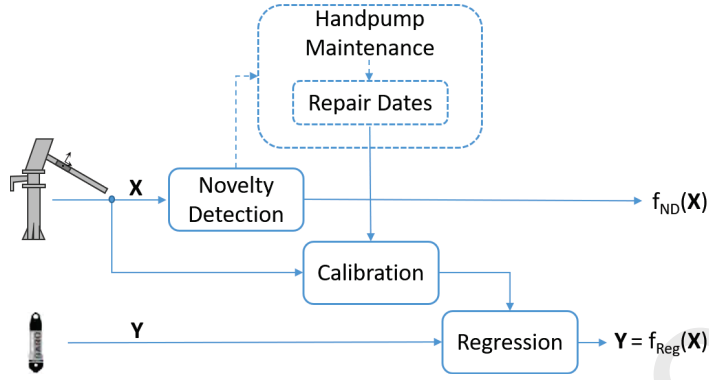


Figure 1: The shallow aquifer monitoring framework designed to work alongside an existing handpump maintenance framework. Vibration data from handpump is used to learn a novelty detection model. Separately, the vibration data along with the water column data measured in the borehole of the handpump are used to learn a regression model. If the handpump is repaired, the vibration data is calibrated before feeding to the regression model.

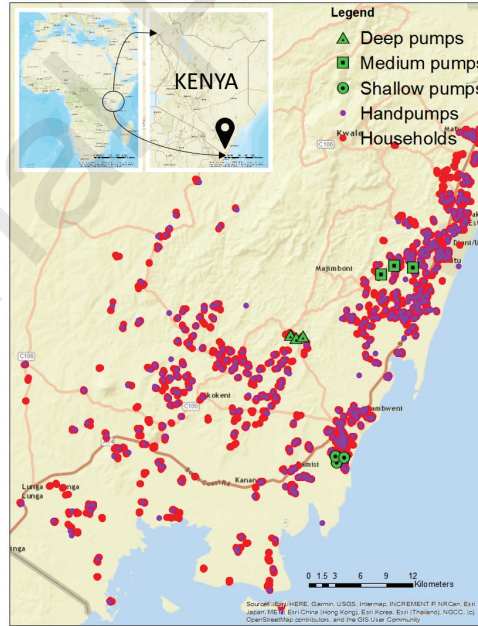


Figure 2: Study area showing households, handpumps, and the sampled handpumps.

characteristics were observed to be different), and the recovery, by definition, is associated with pump disuse, which does not generate any vibration data. The majority of the vibration data collected during pumping corresponds to the available daily water level at the borehole, which in the long-term changes similarly to the daily maximum water level. We exploit this long-term change in the vibration data to model the long-term change in water level at the borehole, which we expect to approximately correspond to the variation in the corresponding shallow aquifer from which the water is being drawn out. We expect the water column variation in the borehole to approximately correspond to the variation in the corresponding shallow aquifer system from which the water is being drawn out.

We use a regression model to learn a mapping function from vibration data to water column. Since we expect vibration data at each handpump to be unique, an independent model is learned for each handpump. To learn such a model, a collection of vibration data and corresponding water column observations are required. These data are collected at three different sites (three handpumps per site) in the study area to test if the model generalizes to different depths of shallow aquifer systems.

A schematic of the shallow aquifer monitoring framework is shown in Fig. 1. The proposed framework is designed to work alongside a handpump maintenance system, represented by dashed lines and not part of the framework itself. The handpump maintenance system monitors the condition of handpumps by tracking the changes in characteristics of vibration data using a novelty detection approach. The outputs of this novelty detection approach, known as novelty scores, may be used to flag a potential pump malfunction

142 if the future vibration data looks substantially different from the previously
 143 observed data, thus facilitating rapid maintenance (Greeff et al., 2019).

144 The proposed technology is based on vibration data obtained from com-
 145 munity handpumps, which are regularly used, and tend to break down quite
 146 frequently, once every few months on average. Depending on the severity of
 147 malfunction and the type of subsequent repair, the characteristics of vibra-
 148 tion data may change substantially, affecting the outputs of the regression
 149 model. The vibration data will also change when the water in the bore-
 150 hole reaches previously unobserved levels. When the characteristics of the
 151 vibration data changes substantially, the regression model learned using pre-
 152 viously observed vibration data may not provide correct estimation. During
 153 such circumstances, the novelty scores may also serve as a guideline to indi-
 154 cate the confidence in the regression model's outputs, where higher novelty
 155 scores would correspond to lower confidence in the regression model's out-
 156 puts.

157 When the vibration data has changed due to a pump malfunction and
 158 subsequent repair, a simple solution to continue using the same regression
 159 model may be to calibrate the post-repair vibration data back to the pre-
 160 repair data. We propose using a regression model to learn a map from post-
 161 repair to pre-repair vibration data by assuming the vibration data averaged
 162 over few days pre vs. post repair are the same. With respect to the water
 163 column estimation model, this simple calibration technique inherently as-
 164 sumes the daily maximum height of water column in the borehole does not
 165 change substantially in a few days, which is generally not a bad assumption
 166 based on Fig.7. To avoid confusing water column-related regression with

167 calibration-related regression, the latter is referred to as calibration in rest
168 of the document.

169 Data from nine handpumps are collected for over a year, out of which two-
170 thirds are used for learning the model (i.e. training for novelty detection, and
171 training and validation for regression), and the remaining one-thirds are used
172 for evaluating the model (i.e. testing).

173 *2.1. Study Area and Sample Selection*

174 The study area is located in Kwale County, Kenya, south of Mombasa and
175 adjacent to northern Tanzania, as shown in Figure 2. From western Shimba
176 Hills to eastern Coastal Plain, the area slopes towards the Indian ocean. The
177 County population of 880,000 people mostly live in rural areas (82%) with
178 the majority (70%) living below the poverty line of less than USD 1.25 a
179 day (KNBS, 2005-06). The study area includes the long-established coastal
180 tourism industry in Diani and the more recent mining and commercial sugar
181 production industries. To sustainably manage the resulting competition for
182 water resources, reliable data on groundwater is vital.

183 A hydrogeological study of the area (Ferrer et al., 2019) reports on the
184 groundwater level change in the coastal coral and sand formations (Pleis-
185 tocene, Magarini, Kalindini) but not the deeper Mazeras sandstone during a
186 La Nina year. They find groundwater levels vary from 4 m below groundwater
187 level (bgl) to 27 m bgl in response to the wetter and drier periods. They also
188 note communities depend on groundwater from boreholes with handpumps
189 as well as open wells in the shallow aquifer. In contrast, the commercial use
190 for irrigation and mining is largely in the deeper aquifer system with limited
191 interaction with the shallow aquifer. The monitoring sites selected for this

Table 1: A summary of handpumps selected at three different sites in the study area. Fig.6a explains the notations.

Depth	Pump Code	Table of Column	Water Column
		TOC (m)	WC (m)
Shallow	SP1	6.4	1.1
	SP2	5.4	1.7
	SP3	8.9	2.0
Medium	MP1	22.2	5.8
	MP2	24.2	2.3
	MP3	17.4	19.4
Deep	DP1	30.4	11.9
	DP2	42.5	6.7
	DP3	35.5	14.9

study are all non-commercial community handpumps. Readers interested in further hydrogeological or socio-economic details of the study area may refer to (Ferrer et al., 2019; Katuva et al., 2019).

To test if the model generalizes to handpumps drawing water from different depths of shallow aquifer systems, three different monitoring sites are selected corresponding to three different depth ranges - shallow, medium, and deep, where these categories are arbitrarily defined based on available samples. Table 1 summarizes table of columns (TOCs) and water columns (WCs) at the boreholes at the time of sensor installation. We also note that all boreholes were mechanically drilled.

The rate of daily drawdown/recovery as well as long-term discharge/recharge

are expected to differ at these three sites due to different hydrogeological, socio-economic, and other factors. Besides, the water columns in the boreholes in the coast are directly affected by both daily and bi-weekly tidal cycles. Pre-processing steps related to removing tidal effects are further described in Section 3.1.

2.2. Regression Approaches for Water Column Estimation

Regression approaches are supervised learning techniques that map examples (inputs) to continuous labels (outputs). We use these approaches to map handpump vibration data (inputs) to their corresponding water column data (outputs). A wide variety of regression approaches exist, e.g. Relevance Vector Machines (RVMs), Gaussian Processes (GPs), decision forest tree regression, multi-layer perceptron regression (MLP regression), etc. Although experiments were performed using different types of regression approaches (yielding comparable results), we focus on MLP (also known as feedforward neural network) regression because they form a natural baseline to compare with Long Short Term Memory (LSTM) networks, the primary approach implemented in this work. Since we wish to provide temporal context to model water column data in terms of past examples of daily handpump vibration data, LSTMs constitute a suitable model for our application.

2.2.1. Recurrent Neural Networks (RNNs)

Feedforward networks do not have feedback, i.e. outputs of the model are not fed back into itself. When feedforward networks are extended to include feedback connections, they are called Recurrent Neural Networks (RNNs). Since feedforward networks assume samples are independent, they are unable




Figure 3: A folded computational graph of RNN (left), and the corresponding unfolded graph (right).




Figure 4: An LSTM cell.

to utilize information from past samples. Recurrent Neural Networks (RNNs) overcome this limitation by allowing information from past samples to persist by using feedback connections. They do so by repeating the feedback loop structure recursively to model a sequence of samples as shown in Fig. 3. The recursive framework allows parameters to be shared across the repetitive structures, which in turn makes parameter estimation feasible. However, although theoretically possible, in practice, RNNs are unable to model long-term dependence because gradients propagated over many stages tend to either vanish (more common) or explode (rarely, but causing much damage to optimization) (Goodfellow et al., 2016).

237 2.2.2. Long Short Term Memory (LSTM) networks

To solve RNN's long-term dependence problem gated RNNs are the most effective solutions designed to date. They do so by allowing gradients to flow for a long duration using gated self-loops. Examples of gated RNNs include Long Short Term Memory (LSTM) networks and gated recurrent units (GRUs) (Goodfellow et al., 2016), the former being the most widely used. LSTMs have the same basic repetitive structure like RNNs, except the repeating structure (called LSTM cell) has three gated self-loops (Fig. 4). A forget gate decides what information will be discarded from the previous cell state. An input gate controls what new information will be stored in the current cell state. Then the outputs from the forget gate and the input gate are used to update the old cell state into a new cell state. Finally, an output gate decides what information will be output from the new cell state. The corresponding equations are

$$f_t = \sigma(\mathbf{W}_f[\mathbf{h}_{t-1}, \mathbf{x}_t] + b_f) \quad (1)$$

$$i_t = \sigma(\mathbf{W}_i[\mathbf{h}_{t-1}, \mathbf{x}_t] + b_i) \quad (2)$$

$$\tilde{C}_t = \tanh(\mathbf{W}_c[\mathbf{h}_{t-1}, \mathbf{x}_t] + b_c) \quad (3)$$

$$C_t = f_t * C_{t-1} + i_t * \tilde{C}_t \quad (4)$$

$$o_t = \sigma(\mathbf{W}_o[\mathbf{h}_{t-1}, \mathbf{x}_t] + b_o) \quad (5)$$



Figure 5: A sequential Neural Network architecture using LSTM.

$$\mathbf{h}_t = o_t * \tanh(C_t), \quad (6)$$

where 1, 2-4, and 5-6 correspond to respectively forget, input, and output gates. Here, C_t is the current cell state, \mathbf{x}_t is the current input vector, \mathbf{h}_t is the current hidden layer vector containing the outputs of all LSTM cells, and b 's and \mathbf{W} 's are respectively biases and weights.

Different variations of LSTMs have been successfully implemented in many applications, e.g. unconstrained handwriting recognition (Graves et al., 2009), speech recognition (Graves, 2013), handwriting generation (Graves, 2013), machine translation (Sutskever et al., 2014), image captioning (Kiros et al., 2014; Vinyals et al., 2014). RNNs and LSTMs have also been implemented to estimate aquifer using hydro-climatic data (Zhang et al., 2018).

Given the relatively small size of the training data, we opt for a simple neural network architecture, consisting of a LSTM unit, a drop-out unit (to prevent over-fitting), and a dense layer in sequence (Fig. 5). As more data becomes available in future, there are opportunities to implement deeper (e.g. stacked layers) and other variations (e.g. shared layers) of networks, which are further discussed in Section 5. The model was implemented in Keras using tensorflow backend. The model parameters (learning rate $[10^{-2} - 10^{-5}]$, number of hidden nodes $[50 - 200]$, epochs $[10 - 200]$, batch size $[10 - 50]$, and input time steps $[1 - 14 \text{ days}]$) were coarsely optimized using separate training and validation sets with (80-20% splits). A separate

test set (one-third of data) was held out for evaluating the trained model. The training, validation, and test sets were split sequentially in time (from past to future) to reflect a real implementation scenario.

2.3. Novelty Detection Approaches for Condition Monitoring

In a standard binary supervised classification approach (Bishop, 2006), models of “normality” and “abnormality” are learned, either separately or simultaneously, using examples of “normal” and “abnormal” observations respectively. When there are not enough examples of abnormal observations, one way to classify these examples is to only learn a model of normality from normal examples, and use this model to classify new examples. Novelty detection is a type of such one-class classification approach, which has been applied in many fields, including electronic IT security, healthcare informatics, industrial monitoring, etc. (Pimentel et al., 2014). This approach is very applicable to condition monitoring in handpumps because usually compared to “normal” handpump operations, there are very limited examples of “abnormal” handpump operations (e.g. broken seal, valve or handle malfunction, etc.). Different types of novelty detection approaches exist, e.g. k nearest neighbour (k -NN), kernel density estimation (KDE), mixture models, one-class Support Vector Machines (one-class SVMs) (Pimentel et al., 2014; Bishop, 2006).

In this work, we use a Gaussian Mixture Model (GMM) (Bishop, 2006) to learn a normal model by representing normal examples with a mixture of multivariate normal densities. Here, an appropriate number of the mixture components is determined based on the data by using Dirichlet process mixtures (Blei and Jordan, 2006) based on a publicly available toolbox (Morton

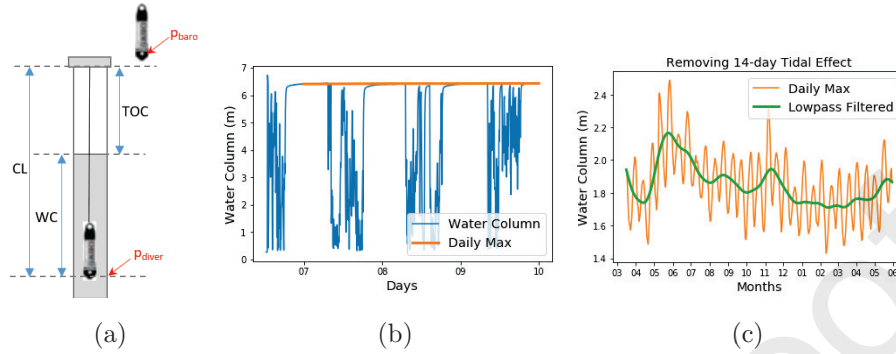


Figure 6: (a) Diver sensor installation close to the bottom of the rising main in the borehole, where the notations mean cable length, table of column, water column, pressure under water and on the surface (Figure adapted from (VanEssen, 2018)), (b-c) Typical variation in water column data in one of the boreholes at (b) the medium depth site and the calculated daily maximum, and (c) the shallow depth site showing daily maximum (that removes daily tidal cycle), and its corresponding low pass filtered signal (that removes bi-weekly tidal cycle).

and Torriane, 2013). For each handpump, a Gaussian Mixture Model is used to learn a normal model based on daily average frequency features generated from handpump vibration data. To help parameter estimation in the mixture model, instead of using all 20 frequency bins, only the first three principal components are considered, which represent more than 95% of total variance in data. The mixture model can be used to compute log-likelihood of examples, the inverse of which can be considered as novelty scores, i.e. lower the log-likelihood, more novel the examples. Since GMM provides probabilistic novelty scores, it is suitable for our application where we intend to use the novelty scores as a measure of confidence of the regression model's outputs.

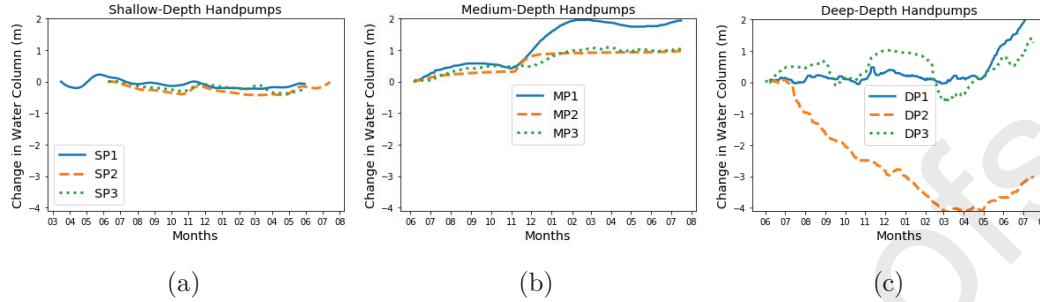


Figure 7: With reference to the start of monitoring date, change in water column data in the boreholes of three sites - (a) Shallow, (b) Medium, and (c) Deep.

3. Data

3.1. Water Column Data

To both learn and evaluate the regression model, water column in the borehole represents the ground truth. The water column is measured using a diver sensor (manufactured by ©Van Essen Instruments) that is fixed close to the bottom of the rising main. Fig. 6a details the nomenclatures used in this document, where water column represents the ground truth for learning the regression model. Data is collected at a sampling interval of five minutes. An example data collected over few days is shown in Fig. 6b, where the blue line shows a typical daily cycle of draw-down during the day, followed by recovery during the night. Since our primary goal is to track the variation in the aquifer level, we use the daily recovered level, approximated by the daily maximum level, as the ground truth for modelling.

Water column in the boreholes in the coast is affected by both daily and bi-weekly tidal cycles. Considering the daily maximum level removes the daily cycle (blue line in Fig. 6c), and low-pass filtering the resulting data

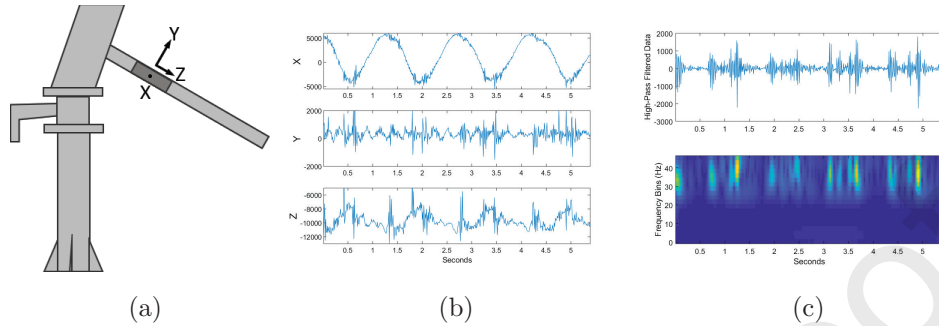


Figure 8: (a) Accelerometer sensor installation close to the fulcrum of the handpump's handle, showing three different acceleration axes, (b) typical vibration data measured at three different axes, and (c) high pass filtered signal of the z-axis vibration data (top) and the corresponding magnitude of its morelet transform (bottom).

removes the bi-weekly cycle (orange line in Fig. 6c). Figs. 7a-7c show the change in water column data in the boreholes at three monitoring sites - shallow, medium, and deep with reference to the start date. These data have been compensated with respect to the atmospheric pressure and pre-processed to represent daily recovered level.

Theoretically, a network of these diver sensors would allow us to directly monitor the shallow aquifer without having to do any modelling. However, given the intended scale of the proposed technology (country to sub-continent), such monitoring program would be cost-prohibitively impractical. Here, the costs not only refer to the cost of the sensor, but also the ones involved in installing the sensor and downloading the data periodically, both extremely laborious process.

321 3.2. Handpump Vibration Data

322 The vibration data is collected using a consumer grade accelerometer
 323 sensor attached to the handle of the handpump as shown in Fig. 8a. The
 324 sensor measures accelerometer data in three orthogonal axes at a sampling
 325 frequency of 95 Hz. A typical pumping cycle, i.e. a cycle of the handle being
 326 raised and pushed down, corresponds to one second interval. An example
 327 data collected over few pumping cycles is shown in Fig. 8b.

328 In machine learning approaches, raw data are often first summarized by
 329 useful information (known as features) before using them to learn a model.
 330 The low frequency component of the data, which corresponds to the pumping
 331 motion, is found to be uninformative of water column variation. The pump-
 332 ing motion is removed using a high-pass filter, and the resulting high-pass
 333 filtered time-series signal (top plot in Fig. 8c) is transformed to frequency
 334 domain using a morelet transform. The magntiude of the morelet transfrom
 335 (bottom plot in Fig. 8c), averaged across the time window, represents the
 336 frequency features corresponding to that window. A daily average of these
 337 temporal windows along with the corresponding daily maximum water col-
 338 umn represent a feature-label pair. A collection of these feature-label pairs
 339 per handpump constitutes a training and testing dataset for that handpump.

340 Vibration data maybe unavailable for a period of time due to several
 341 reasons, most commonly due to battery or sensor malfunction, or handpump
 342 disuse. A handpump may be temporarily out of operation usually because
 343 of malfunction or alternate water sources (e.g. surface water) being available
 344 during rainy season, or school closures in case of handpumps installed at
 345 the schools. We propose using Gaussian Processes (Rasmussen, 2004) to

impute frequency features corresponding to missing days. Unfortunately, for two handpumps, one shallow and one deep, too much data were missing to allow reasonable data imputation. Hence, we only report results on the seven remaining handpumps. Since the highest proportion of data is available for the medium depth handpumps, we provide detailed results for that site while only providing key results for the other two sites.

4. Results

The machine learning approaches described in Section 2 are used to perform several experiments related to water column estimation using the dataset described in Section 3. At this stage, our primary goal is to assess the feasibility of the proposed technology in a real world setting, i.e. whether vibration data from a collection of community handpumps at a location can be used to infer the trend in water column variation of the corresponding aquifer.

4.1. Novelty Detection Analysis

Figs. 9-11 show the novelty scores of training (blue dots) and test (red dots) examples in terms of log-likelihood of examples given the normal model. In general, the examples in future incrementally appear to be more different from the normal examples. This trend is expected because the vibration data is expected to change over time either due to gradual pump wear and/or change in water level in the borehole to previously unobserved levels (a relatively smooth change), or due to severe pump malfunction and subsequent repair (a relatively abrupt change). A continuous drop in the log-likelihood, suggesting a gradual change in vibration data, may correspond to change in

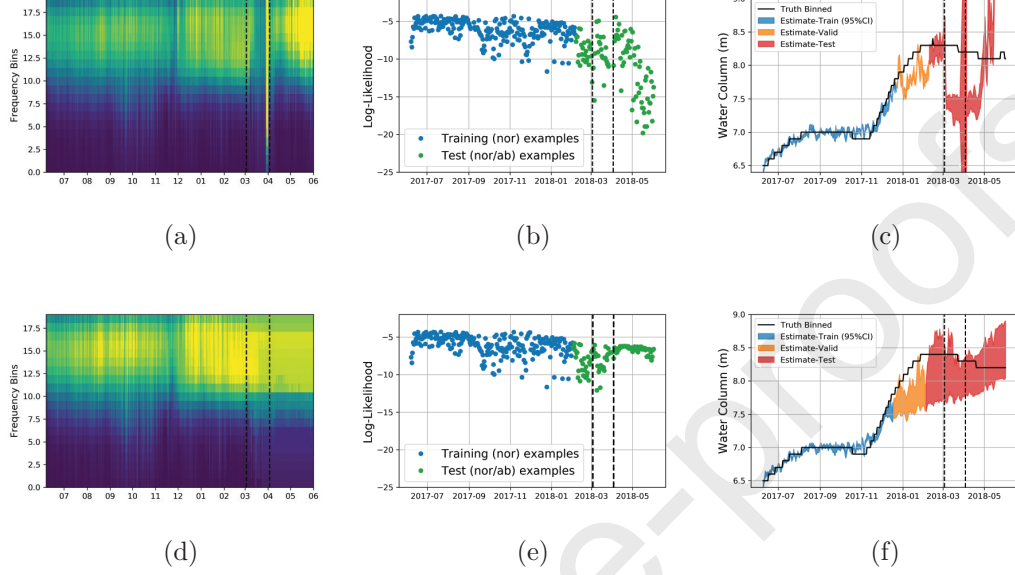


Figure 9: For a particular handpump (MP1), (a) features generated from vibration data, (b) log-likelihood of examples given normal model, and (c) water column estimates using LSTM. These plots are redrawn in (d)-(f) once the features are calibrated at handpump repair dates, represented by the dashed lines.

370 water column over time (e.g. Fig. 10c). On the other hand, an abrupt change
 371 in vibration data due to severe pump malfunction/repair often corresponds
 372 to a sharp drop in the log-likelihood. Most of these cases (e.g. Figs. 10a,
 373 10b, etc.) stand out visually, and are aligned to their corresponding pump
 374 repair dates (black dashed lines), whenever the repair dates are available.
 375 These insights from novelty detection analysis help understand the results of
 376 experiments related to calibration and water column estimation.

377 4.2. *Vibration Data Calibration*

378 Usually the novelty scores appear to be higher (i.e. the log-likelihood
 379 values are lower) immediately following a repair (e.g. Fig. 9b), indicating a
 380 substantial change in the vibration data after the repair. To enable using a
 381 single regression model to estimate water column, the post-repair vibration
 382 data is calibrated to match the pre-repair data. An example of this cali-
 383 bration process is detailed in Fig. 9, where the top row shows the features
 384 generated from the vibration data (Fig. 9a), the log-likelihoods (Fig. 9b),
 385 and the water column estimates (Fig. 9c), and the bottom row shows the
 386 corresponding plots after calibration. Since a very simple regression model
 387 was used to calibrate the vibration data, its quality is expected to degrade
 388 as further calibrations are performed in future. In the example, the second
 389 calibration does not work as well as the first calibration. As part of our fu-
 390 ture work, we intend to explore more principled ways to model this problem,
 391 e.g. transfer learning, which is discussed in Section 5.

392 4.3. *Water Column Estimation Analysis*

393 For medium-depth handpumps, results are shown for LSTM, and for ref-
 394 erence, compared to MLP regression in Fig. 10. For shallow and deep hand-
 395 pumps, results are only shown for LSTM in Fig. 11. In both figures, each
 396 column corresponds to results for a particular handpump, where the rows
 397 correspond to the novelty scores, novelty scores post-calibration, LSTM es-
 398 timates, and MLP estimates (only in Fig. 10) from top to bottom. The
 399 estimates for training, validation, and test sets are shown in blue, orange,
 400 and red colours respectively. The estimates are shown in terms of their
 401 95% confidence interval based on 10 iterations of training LSTM model with

Table 2: For each regressor, mean error (metres) on test dataset (top row), and one standard deviation (bottom row).

	SP1	SP2	MP1	MP2	MP3	DP1	DP2
MLP	0.14	0.26	0.31	0.29	0.18	0.60	0.85
	0.12	0.19	0.15	0.11	0.10	0.34	0.71
LSTM	0.09	0.12	0.23	0.14	0.15	0.43	1.07
	0.04	0.06	0.17	0.11	0.12	0.47	0.41

random initialization. While parameters for each approach were coarsely optimized using cross-validation, we have not performed exhaustive parameter optimization because our primary goal at this stage is to assess whether the estimates roughly tracks the true changes in water column rather than obtain the best possible outputs. A summary of results for all handpumps for both MLP and LSTM techniques are provided in Table 2.

Results show that, in general, vibration data are indicative of changes in the water column. LSTM generally outperforms standard MLP regression technique, where LSTM is expected to perform even better as more training data become available. In most cases, pump repairs change the vibration data features substantially, which when “corrected” via calibration, does somewhat help to improve water column estimation. But the frequency and/or the nature of repairs may determine the efficacy of the simple calibration technique implemented in this work. Nevertheless, the novelty detection outputs provide a reasonably accurate guideline to determine when to trust the regression model outputs. Typically a drop in the log-likelihood corresponds to either an inaccurate water column estimate or one with high

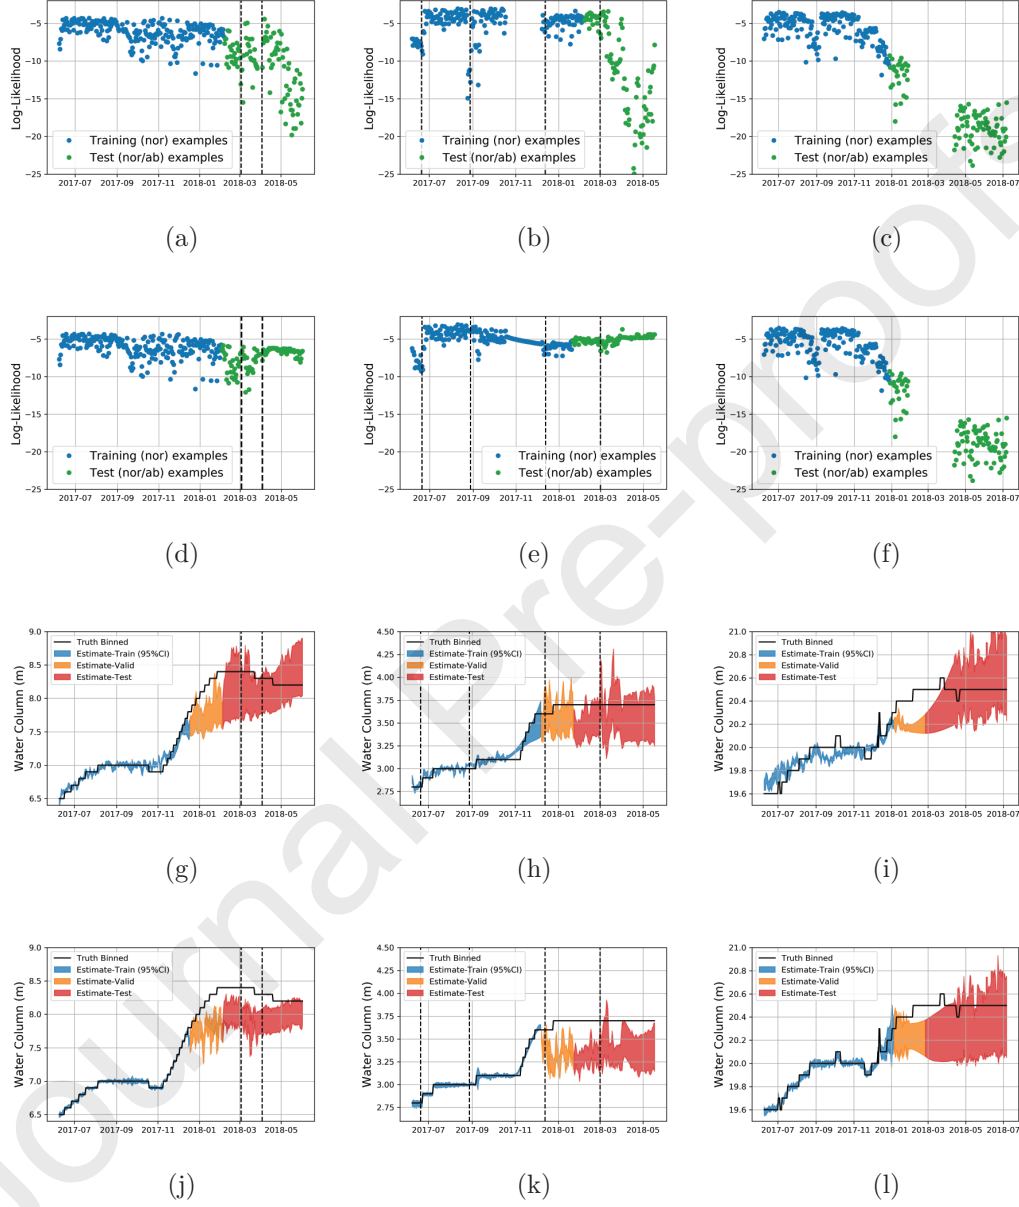


Figure 10: For each column, from top-bottom rows, log-likelihood of examples given normal model, log-likelihood after calibration, water column estimates using LSTM and MLP regressor. Columns correspond to medium-depth handpumps MP1, MP2, and MP3 in the Ukunda area.

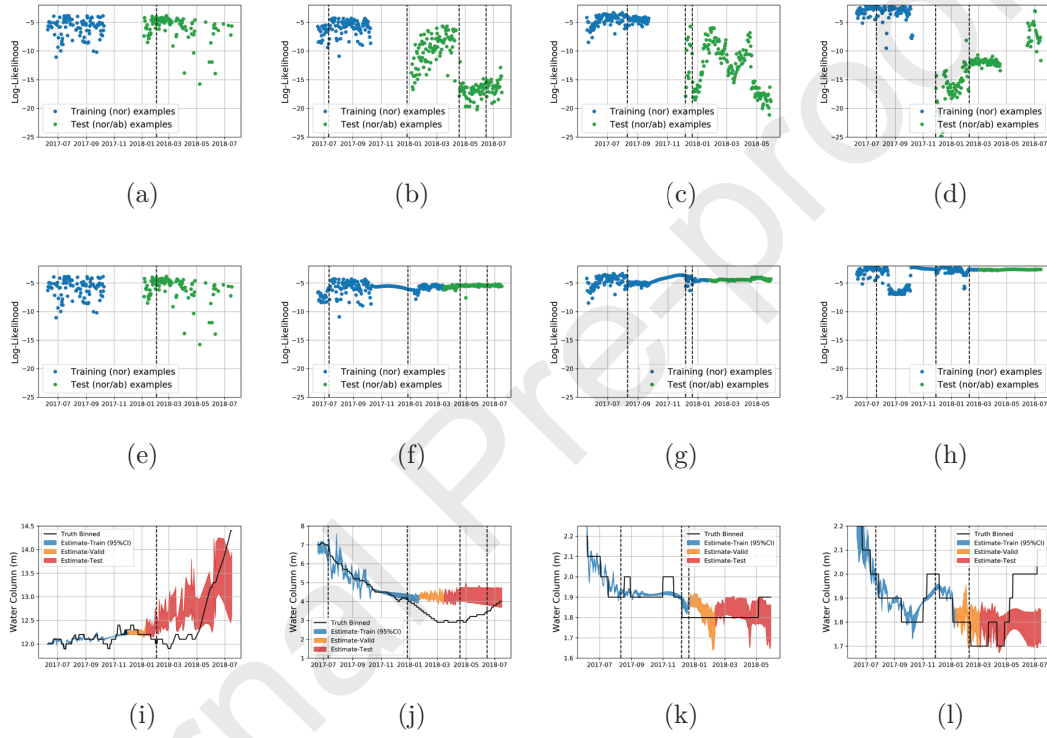


Figure 11: For each column, from top-bottom rows, log-likelihood of examples given normal model, log-likelihood after calibration, and water column estimates using LSTM. Columns correspond to deep handpumps DP1, DP2 and shallow handpumps SP1, SP2 at the Shimba Hills and the Coast sites respectively.

419 uncertainty.

420 The proposed technology is intended to be implemented at scale by con-
 421 currently modeling a network of community handpumps. Rather than es-
 422 timating continuous changes in water column at individual boreholes, our
 423 ultimate goal is to use a network of handpumps to infer the overall trend
 424 in the water levels in the corresponding shallow aquifer. With this goal in
 425 mind, we plot the fractional change in water columns at the boreholes at
 426 two locations with respect to a common reference date. Figs. 12a and 12b
 427 show these plots for medium and shallow handpumps respectively. These
 428 results show the estimated changes in water column approximately track the
 429 true trend. However, the estimates deteriorate as we start predicting further
 430 ahead in time due to the limitations in the model, which are discussed in
 431 Section 4.4. Similar plots are not shown for deep handpumps because the
 432 difference in their water column data variation suggests that there may be
 433 differences in the hydrogeological dynamics of the aquifers from which the
 434 two deep handpumps are drawing water.

435 4.4. *Limitations*

436 Despite the resources invested in collecting a unique handpump vibration
 437 dataset using novel sensing technology, in terms of hydrogeology, one year
 438 of data is insufficient to track long-term changes in aquifer level. Few years
 439 of data may be required to capture a typical range of aquifer level changes.
 440 In the current dataset, validation and test data are often very different from
 441 training data, complicating both training and testing the model.

442 Approximately 15% of data were missing, usually over continuous streams
 443 of time (from days to months), primarily due to discharged batteries or other




Figure 12: Fractional change in water column with respect to the reference water column observed on 2017/6/1 for the handpumps in (a) medium-depth and (b) shallow-depth clusters.

444 hardware-related problems, handpump malfunction, or temporary handpump
 445 disuse because of availability of alternate water sources during rainy season
 446 (e.g. rain harvesting, surface water). Despite using Gaussian Processes to
 447 impute the missing data, the accuracy of the imputed data degrades as the
 448 duration of missing data increases.

449 The effectiveness of the vibration data calibration technique depends on
 450 the frequency and type of handpump repair, with the calibration being less
 451 effective as the number and/or severity of breakdowns/repairs increase. Al-
 452 though the novelty scores provide a guideline to assess the confidence of the
 453 water column estimates, better techniques are required to model this problem
 454 in a principled manner, e.g. transfer learning.

5. Conclusions

The proposed technology based on handpump vibration data and machine learning approaches show promise in providing useful aquifer monitoring information in resource-constrained regions where alternate sources of data are either scarce or cost-prohibitive and where handpumps already constitute part of the rural water infrastructure. In regions where handpump maintenance infrastructure is already in place, there is an opportunity to use the vibration data for the additional purpose of providing rudimentary shallow aquifer monitoring at minimal extra cost. As expected, going from constrained to unconstrained real-world environment brings challenges. The most pressing one concerns frequent handpump breakdown and subsequent repair that complicates learning a consistent model to estimate changes in water level. Beside the solutions implemented in this work in the form of novelty scores and vibration data calibration, we intend to explore more principled methods in future to model this problem. Given a pre-trained model, transfer learning (Zoph et al., 2016; Yang et al., 2017) allows learning a variation of the learned model for a dataset with different characteristics by only requiring a relatively smaller dataset. Further field-experiments will be required to determine if transfer learning is feasible, and how much new data (hours-days) will be required to re-train the model. We note that in practice, the proposed technology is designed to be implemented at scale using a network of handpumps. As long as a sufficiently large number of these handpumps at a location are operational, the framework may allow useful approximate inference about the underlying aquifer behaviour within a degree of uncertainty. Further large-scale field experiments are required to ascertain

the viability of this technology for real-world application as well as to determine under which conditions this technology is likely to fail and succeed. We note that the proposed technology is aimed to complement cost-effectively any traditional hydrogeological models rather than replace them.

As part of our future work, there are opportunities to model multiple handpumps simultaneously and fuse hydro-climatic data, and wherever available, outputs from hydrogeological models. Multi-task learning extensions of LSTM (Liu et al., 2016; Chen et al., 2018) are well-suited to model such scenarios.

Given the increasing global importance of groundwater monitoring technology, there needs to be more research to fill the gap between available state-of-the-art but cost-prohibitive technologies and the capacity of developing nations to adopt them. Novel cost-effective technologies that utilize the available infrastructure of these regions may be a plausible path moving forward. Given the success and ongoing trials of machine learning approaches to solve global socio-developmental challenges (Beckel et al., 2013; Blumenstock et al., 2015; Jean et al., 2016; Rolnick et al., 2019), experimentation of these approaches may provide cost-effective complementary options to boost on-going progress in hydrogeological modeling to monitor groundwater.

Acknowledgment

The authors would like to thank the Kenyan team (Nanga, Dola, George, Sidney, Rueben, James Okoti, Idd Mwaropia, Suleiman Mwakurya, Fauzia Swaleh), FundiFix, Rural Focus Ltd., Base Titanium Ltd., and the Kwale Country Government. This research was funded by the UK Government

via NERC, ESRC, and DFID as part of the Gro for GooD project (UPGro Consortium Grant: NE/M008894/1).

Achieng, K.O., 2019. Modelling available water capacity of topsoil in a bayesian paradigm. *Environmental Modelling Software* 120, 104500. URL: <http://www.sciencedirect.com/science/article/pii/S1364815219304165>, doi:<https://doi.org/10.1016/j.envsoft.2019.104500>.

Beckel, C., Sadamori, L., Santini, S., 2013. Automatic socio-economic classification of households using electricity consumption data, in: *Proceedings of the Fourth International Conference on Future Energy Systems*, ACM, New York, NY, USA. pp. 75–86. URL: <http://doi.acm.org/10.1145/2487166.2487175>, doi:10.1145/2487166.2487175.

Behzad, M., Asghari, K., Jr., E.A.C., 2010. Comparative study of svms and anns in aquifer water level prediction. *Journal of Computing in Civil Engineering* 24. doi:10.1061/ASCECP.1943-5487.0000043.

Bishop, C.M., 2006. *Pattern Recognition and Machine Learning (Information Science and Statistics)*. Springer-Verlag, Berlin, Heidelberg.

Blei, D.M., Jordan, M.I., 2006. Variational inference for dirichlet process mixtures. *Bayesian Anal.* 1, 121–143. URL: <https://doi.org/10.1214/06-BA104>, doi:10.1214/06-BA104.

Blumenstock, J., Cadamuro, G., On, R., 2015. Predicting poverty and wealth from mobile phone metadata. *Science* 350, 1073–1076. URL: <http://science.sciencemag>.

- org/content/350/6264/1073, doi:10.1126/science.aac4420,
arXiv:<http://science.sciencemag.org/content/350/6264/1073.full.pdf>.
- Chen, J., Qiu, X., Liu, P., Huang, X., 2018. Meta Multi-Task Learning for Sequence Modeling. arXiv e-prints arXiv:1802.08969.
- Colchester, F.E., Greeff, H., Thomson, P., Hope, R., Clifton, D.A., 2014. Smart handpumps: a preliminary data analysis. IET Conference Proceedings , 7–7(1)URL: <https://digital-library.theiet.org/content/conferences/10.1049/cp.2014.0767>.
- Colchester, F.E., Marais, H.G., Thomson, P., Hope, R., Clifton, D.A., 2017. Accidental infrastructure for groundwater monitoring in africa. Environmental Modelling Software 91, 241 – 250. URL: <http://www.sciencedirect.com/science/article/pii/S1364815216308325>, doi:<https://doi.org/10.1016/j.envsoft.2017.01.026>.
- Comte, J.C., Cassidy, R., Obando, J., Robins, N., Ibrahim, K., Melchiorly, S., Mjemah, I., Shauri, H., Bourhane, A., Mohamed, I., Noe, C., Mwega, B., Makokha, M., Join, J.L., Banton, O., Davies, J., 2016. Challenges in groundwater resource management in coastal aquifers of east africa: Investigations and lessons learnt in the comoros islands, kenya and tanzania. Journal of Hydrology: Regional Studies 5, 179 – 199. URL: <http://www.sciencedirect.com/science/article/pii/S2214581815002232>, doi:<https://doi.org/10.1016/j.ejrh.2015.12.065>.

- 550 Fan, Y., Li, H., Miguez-Macho, G., 2013. Global patterns of groundwater
551 table depth. *Science* 339, 940–943. URL: [https://science.sciencemag.](https://science.sciencemag.org/content/339/6122/940)
552 [org/content/339/6122/940](https://science.sciencemag.org/content/339/6122/940), doi:10.1126/science.1229881,
553 arXiv:<https://science.sciencemag.org/content/339/6122/940.full.pdf>.
- 554 FAO, 2011. The state of the world's land and water resources for food and
555 agriculture (solaw) – managing systems at risk.
- 556 Ferrer, N., Folch, A., Lane, M., Olago, D., Odida, J., Custodio, E., 2019.
557 Groundwater hydrodynamics of an eastern africa coastal aquifer, includ-
558 ing la niÃ±a 2016–17 drought. *Science of The Total Environment* 661,
559 575 – 597. URL: [http://www.sciencedirect.com/science/article/](http://www.sciencedirect.com/science/article/pii/S0048969719302177)
560 [pii/S0048969719302177](http://www.sciencedirect.com/science/article/pii/S0048969719302177), doi:[https://doi.org/10.1016/j.scitotenv.](https://doi.org/10.1016/j.scitotenv.2019.01.198)
561 [2019.01.198](https://doi.org/10.1016/j.scitotenv.2019.01.198).
- 562 Goodfellow, I., Bengio, Y., Courville, A., 2016. *Deep Learning*. MIT Press.
563 <http://www.deeplearningbook.org>.
- 564 Graves, A., 2013. Generating sequences with recurrent neural net-
565 works. CoRR abs/1308.0850. URL: <http://arxiv.org/abs/1308.0850>,
566 arXiv:1308.0850.
- 567 Graves, A., Liwicki, M., Fernández, S., Bertolami, R., Bunke, H., Schmidhu-
568 ber, J., 2009. A novel connectionist system for unconstrained handwriting
569 recognition. *IEEE Transactions on Pattern Analysis and Machine Intelli-*
570 *gence* 31, 855–868. doi:10.1109/TPAMI.2008.137.
- 571 Greeff, H., Manandhar, A., Thomson, P., 2020. Daily handpump
572 accelerometer data and borehole water level data, kwale county,

- kenya (dataset). British Geological Survey URL: <https://dx.doi.org/10.5285/2a7ed1a6-4749-457b-86a3-e876d849338d>, doi:10.5285/2a7ed1a6-4749-457b-86a3-e876d849338d.
- Greeff, H., Manandhar, A., Thomson, P., Hope, R., Clifton, D.A., 2019. Distributed inference condition monitoring system for rural infrastructure in the developing world. *IEEE Sensors Journal* 19, 1820–1828. doi:10.1109/JSEN.2018.2882866.
- Herrera, M., Torgo, L., Izquierdo, J., Pérez-García, R., 2010. Predictive models for forecasting hourly urban water demand. *Journal of Hydrology* 387, 141 – 150. URL: <http://www.sciencedirect.com/science/article/pii/S0022169410001861>, doi:<https://doi.org/10.1016/j.jhydrol.2010.04.005>.
- Hope, R., Rouse, M., 2013. Risks and responses to universal drinking water security. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 371, 20120417. URL: <https://royalsocietypublishing.org/doi/abs/10.1098/rsta.2012.0417>, doi:10.1098/rsta.2012.0417, arXiv:<https://royalsocietypublishing.org/doi/pdf/10.1098/rsta.2012.0417>.
- Howard, K., 2014. Sustainable cities and the groundwater governance challenge. *Environmental Earth Sciences* 73, 2543–2554. doi:10.1007/s12665-014-3370-y.
- Jean, N., Burke, M., Xie, M., Davis, W.M., Lobell, D.B., Ermon, S., 2016. Combining satellite imagery and machine learning to predict

- poverty. *Science* 353, 790–794. URL: <http://science.sciencemag.org/content/353/6301/790>, doi:10.1126/science.aaf7894, arXiv:<http://science.sciencemag.org/content/353/6301/790.full.pdf>.
- Katuva, J., Hope, R., Foster, T., Koehler, J., Thomson, P., 2019. Groundwater and welfare: A conceptual framework applied to coastal kenya. *Journal of Groundwater for Sustainable Development* (Accepted) .
- Kelbe, B.E., Grundling, A.T., Price, J.S., 2016. Modelling water-table depth in a primary aquifer to identify potential wetland hydrogeomorphic settings on the northern maputaland coastal plain, kwazulu-natal, south africa. *Hydrogeology Journal* 24, 249–265. URL: <https://doi.org/10.1007/s10040-015-1350-2>, doi:10.1007/s10040-015-1350-2.
- Kiros, R., Salakhutdinov, R., Zemel, R.S., 2014. Unifying visual-semantic embeddings with multimodal neural language models. *CoRR* abs/1411.2539. URL: <http://arxiv.org/abs/1411.2539>, arXiv:1411.2539.
- KNBS, 2005-06. Kenya integrated household budget survey, government of kenya: Kenya national bureau of statistics .
- Liu, P., Qiu, X., Huang, X., 2016. Recurrent neural network for text classification with multi-task learning. *CoRR* abs/1605.05101. URL: <http://arxiv.org/abs/1605.05101>, arXiv:1605.05101.
- MacDonald, A.M., Bonsor, H.C., Dochartaigh, B.É.Ó., Taylor, R.G., 2012. Quantitative maps of groundwater resources in africa. *Environ-*

- 618 mental Research Letters 7, 024009. URL: [https://doi.org/10.1088/](https://doi.org/10.1088/2F1748-9326%2F7%2F2%2F024009)
619 [2F1748-9326/7/2/024009](https://doi.org/10.1088/1748-9326/7/2/024009), doi:10.1088/1748-9326/7/2/024009.
- 620 Manandhar, A., 2019. Water column modelling using lstm (software). Github
621 URL: <https://github.com/achutman/handpumpAquiferLstm>.
- 622 Morton, K., Torrione, P., 2013. PRT: Pattern recognition and machine learn-
623 ing in matlab (software). Github URL: [https://github.com/covartech/](https://github.com/covartech/PRT)
624 PRT.
- 625 Nayak, P.C., Rao, Y.R.S., Sudheer, K.P., 2006. Groundwater level fore-
626 casting in a shallow aquifer using artificial neural network approach. Wa-
627 ter Resources Management 20, 77–90. URL: [https://doi.org/10.1007/](https://doi.org/10.1007/s11269-006-4007-z)
628 [s11269-006-4007-z](https://doi.org/10.1007/s11269-006-4007-z), doi:10.1007/s11269-006-4007-z.
- 629 Pimentel, M.A.F., Clifton, D.A., Clifton, L.A., Tarassenko, L., 2014. A
630 review of novelty detection. Signal Processing 99, 215–249.
- 631 Rasmussen, C.E., 2004. Gaussian Processes in Machine Learning. Springer
632 Berlin Heidelberg, Berlin, Heidelberg. pp. 63–71. URL: [https://doi.org/](https://doi.org/10.1007/978-3-540-28650-9_4)
633 [10.1007/978-3-540-28650-9_4](https://doi.org/10.1007/978-3-540-28650-9_4), doi:10.1007/978-3-540-28650-9_4.
- 634 Richey, A., Thomas, B., Lo, M.H., Famiglietti, J., Swenson, S., Rodell,
635 M., 2015. Uncertainty in global groundwater storage estimates in a
636 total groundwater stress framework. Water Resources Research 51,
637 5198–5216. URL: [https://www2.scopus.com/inward/record.uri?eid=](https://www2.scopus.com/inward/record.uri?eid=2-s2.0-84939468730&doi=10.1002%2f2015WR017351&partnerID=40&md5=458b45c0f261692831cd01acc2252c54)
638 [2-s2.0-84939468730&doi=10.1002%2f2015WR017351&partnerID=40&](https://www2.scopus.com/inward/record.uri?eid=2-s2.0-84939468730&doi=10.1002%2f2015WR017351&partnerID=40&md5=458b45c0f261692831cd01acc2252c54)
639 [md5=458b45c0f261692831cd01acc2252c54](https://www2.scopus.com/inward/record.uri?eid=2-s2.0-84939468730&doi=10.1002%2f2015WR017351&partnerID=40&md5=458b45c0f261692831cd01acc2252c54), doi:10.1002/2015WR017351.
640 cited By 64.

- 641 Robert, C., 2014. Machine learning, a probabilistic perspec-
 642 tive. CHANCE 27, 62–63. URL: [https://doi.org/10.1080/](https://doi.org/10.1080/09332480.2014.914768)
 643 09332480.2014.914768, doi:10.1080/09332480.2014.914768,
 644 arXiv:<https://doi.org/10.1080/09332480.2014.914768>.
- 645 Rolnick, D., Donti, P.L., Kaack, L.H., Kochanski, K., Lacoste, A., Sankaran,
 646 K., Ross, A.S., Milojevic-Dupont, N., Jaques, N., Waldman-Brown, A.,
 647 Luccioni, A., Maharaj, T., Sherwin, E.D., Mukkavilli, S.K., Kording,
 648 K.P., Gomes, C., Ng, A.Y., Hassabis, D., Platt, J.C., Creutzig, F.,
 649 Chayes, J., Bengio, Y., 2019. Tackling climate change with machine learn-
 650 ing. CoRR abs/1906.05433. URL: <http://arxiv.org/abs/1906.05433>,
 651 arXiv:1906.05433.
- 652 Sahoo, S., Russo, T.A., Elliott, J., Foster, I., 2017. Machine
 653 learning algorithms for modeling groundwater level changes in
 654 agricultural regions of the u.s. Water Resources Research 53,
 655 3878–3895. URL: [https://agupubs.onlinelibrary.wiley.com/](https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1002/2016WR019933)
 656 doi/abs/10.1002/2016WR019933, doi:10.1002/2016WR019933,
 657 arXiv:<https://agupubs.onlinelibrary.wiley.com/doi/pdf/10.1002/2016WR019933>.
- 658 Suryanarayana, C., Sudheer, C., Mahammood, V., Panigrahi, B.,
 659 2014. An integrated wavelet-support vector machine for ground-
 660 water level prediction in visakhapatnam, india. Neurocomput-
 661 ing 145, 324 – 335. URL: [http://www.sciencedirect.com/](http://www.sciencedirect.com/science/article/pii/S0925231214006407)
 662 science/article/pii/S0925231214006407, doi:[https://doi.org/10.](https://doi.org/10.1016/j.neucom.2014.05.026)
 663 1016/j.neucom.2014.05.026.
- 664 Sutskever, I., Vinyals, O., Le, Q.V., 2014. Sequence to se-

- 665 quence learning with neural networks, in: Ghahramani, Z., Welling,
 666 M., Cortes, C., Lawrence, N.D., Weinberger, K.Q. (Eds.), Ad-
 667 vances in Neural Information Processing Systems 27. Curran Asso-
 668 ciates, Inc., pp. 3104–3112. URL: [http://papers.nips.cc/paper/](http://papers.nips.cc/paper/5346-sequence-to-sequence-learning-with-neural-networks.pdf)
 669 [5346-sequence-to-sequence-learning-with-neural-networks.pdf](http://papers.nips.cc/paper/5346-sequence-to-sequence-learning-with-neural-networks.pdf).
- 670 Taormina, R., wing Chau, K., Sethi, R., 2012. Artificial neural network
 671 simulation of hourly groundwater levels in a coastal aquifer system of the
 672 venice lagoon. Engineering Applications of Artificial Intelligence 25, 1670
 673 – 1676. URL: [http://www.sciencedirect.com/science/article/pii/](http://www.sciencedirect.com/science/article/pii/S0952197612000462)
 674 [S0952197612000462](http://www.sciencedirect.com/science/article/pii/S0952197612000462), doi:[https://doi.org/10.1016/j.engappai.2012.](https://doi.org/10.1016/j.engappai.2012.02.009)
 675 [02.009](https://doi.org/10.1016/j.engappai.2012.02.009).
- 676 Tapoglou, E., Karatzas, G.P., Trichakis, I.C., Varouchakis, E.A.,
 677 2014. A spatio-temporal hybrid neural network-kriging model for
 678 groundwater level simulation. Journal of Hydrology 519, 3193 –
 679 3203. URL: [http://www.sciencedirect.com/science/article/pii/](http://www.sciencedirect.com/science/article/pii/S002216941400835X)
 680 [S002216941400835X](http://www.sciencedirect.com/science/article/pii/S002216941400835X), doi:[https://doi.org/10.1016/j.jhydrol.2014.](https://doi.org/10.1016/j.jhydrol.2014.10.040)
 681 [10.040](https://doi.org/10.1016/j.jhydrol.2014.10.040).
- 682 Thomson, P., Hope, R., Foster, T., 2012. GSM-enabled re-
 683 mote monitoring of rural handpumps: a proof-of-concept
 684 study. Journal of Hydroinformatics 14, 829–839. URL: <https://doi.org/10.2166/hydro.2012.183>,
 685 [doi:10.2166/hydro.2012.183](https://doi.org/10.2166/hydro.2012.183),
 686 [arXiv:https://iwaponline.com/jh/article-pdf/14/4/829/386770/829.pdf](https://iwaponline.com/jh/article-pdf/14/4/829/386770/829.pdf).
- 687 UN, 2018. Sustainable development goal 6 synthesis report on water and
 688 sanitation.

- 689 Van Camp, M., Mjemah, I.C., Al Farrah, N., Walraevens, K., 2013.
 690 Modeling approaches and strategies for data-scarce aquifers: example
 691 of the dar es salaam aquifer in tanzania. *Hydrogeology Journal* 21,
 692 341–356. URL: <https://doi.org/10.1007/s10040-012-0908-5>, doi:10.
 693 1007/s10040-012-0908-5.
- 694 VanEssen, 2018. Product manual td-divertm baro-diver®
 695 â di8xx series. [https://www.vanessen.com/images/PDFs/](https://www.vanessen.com/images/PDFs/TD-Diver-DI8xx-ProductManual-en.pdf)
 696 [TD-Diver-DI8xx-ProductManual-en.pdf](https://www.vanessen.com/images/PDFs/TD-Diver-DI8xx-ProductManual-en.pdf).
- 697 Vinyals, O., Toshev, A., Bengio, S., Erhan, D., 2014. Show and tell: A
 698 neural image caption generator. *CoRR abs/1411.4555*. URL: [http://](http://arxiv.org/abs/1411.4555)
 699 arxiv.org/abs/1411.4555, arXiv:1411.4555.
- 700 WWAP, 2019. The united nations world water development report 2019:
 701 Leaving no one behind.
- 702 Xu, X., Huang, G., Zhan, H., Qu, Z., Huang, Q., 2012. Integra-
 703 tion of swap and modflow-2000 for modeling groundwater dynamics
 704 in shallow water table areas. *Journal of Hydrology* 412-413, 170 –
 705 181. URL: [http://www.sciencedirect.com/science/article/pii/](http://www.sciencedirect.com/science/article/pii/S002216941100429X)
 706 [S002216941100429X](http://www.sciencedirect.com/science/article/pii/S002216941100429X), doi:[https://doi.org/10.1016/j.jhydrol.2011.](https://doi.org/10.1016/j.jhydrol.2011.07.002)
 707 [07.002](https://doi.org/10.1016/j.jhydrol.2011.07.002). hydrology Conference 2010.
- 708 Yang, Z., Salakhutdinov, R., Cohen, W.W., 2017. Transfer
 709 learning for sequence tagging with hierarchical recurrent networks.
 710 *CoRR abs/1703.06345*. URL: <http://arxiv.org/abs/1703.06345>,
 711 [arXiv:1703.06345](http://arxiv.org/abs/1703.06345).

- 712 Yaseen, Z.M., Sulaiman, S.O., Deo, R.C., Chau, K.W., 2019. An
 713 enhanced extreme learning machine model for river flow forecasting:
 714 State-of-the-art, practical applications in water resource engineering
 715 area and future research direction. *Journal of Hydrology* 569, 387
 716 – 408. URL: [http://www.sciencedirect.com/science/article/pii/](http://www.sciencedirect.com/science/article/pii/S0022169418309545)
 717 [S0022169418309545](http://www.sciencedirect.com/science/article/pii/S0022169418309545), doi:[https://doi.org/10.1016/j.jhydrol.2018.](https://doi.org/10.1016/j.jhydrol.2018.11.069)
 718 [11.069](https://doi.org/10.1016/j.jhydrol.2018.11.069).
- 719 Yoon, H., Jun, S.C., Hyun, Y., Bae, G.O., Lee, K.K., 2011. A com-
 720 parative study of artificial neural networks and support vector ma-
 721 chines for predicting groundwater levels in a coastal aquifer. *Jour-*
 722 *nal of Hydrology* 396, 128 – 138. URL: [http://www.sciencedirect.](http://www.sciencedirect.com/science/article/pii/S0022169410006761)
 723 [com/science/article/pii/S0022169410006761](http://www.sciencedirect.com/science/article/pii/S0022169410006761), doi:[https://doi.org/](https://doi.org/10.1016/j.jhydrol.2010.11.002)
 724 [10.1016/j.jhydrol.2010.11.002](https://doi.org/10.1016/j.jhydrol.2010.11.002).
- 725 Zhang, J., Zhu, Y., Zhang, X., Ye, M., Yang, J., 2018. Develop-
 726 ing a long short-term memory (lstm) based model for predicting wa-
 727 ter table depth in agricultural areas. *Journal of Hydrology* 561, 918
 728 – 929. URL: [http://www.sciencedirect.com/science/article/pii/](http://www.sciencedirect.com/science/article/pii/S0022169418303184)
 729 [S0022169418303184](http://www.sciencedirect.com/science/article/pii/S0022169418303184), doi:[https://doi.org/10.1016/j.jhydrol.2018.](https://doi.org/10.1016/j.jhydrol.2018.04.065)
 730 [04.065](https://doi.org/10.1016/j.jhydrol.2018.04.065).
- 731 Zoph, B., Yuret, D., May, J., Knight, K., 2016. Transfer learning for low-
 732 resource neural machine translation. *CoRR* abs/1604.02201. URL: [http:](http://arxiv.org/abs/1604.02201)
 733 [//arxiv.org/abs/1604.02201](http://arxiv.org/abs/1604.02201), arXiv:1604.02201.