

A technological framework for data-driven IoT systems: application on landslide monitoring.

Sivanarayani M Karunarathne^α, Matthew Dray^α, Lyudmil Popov^α, Matthew Butler^α,
Catherine Pennington^β, Constantin Marios Angelopoulos^α

^αBournemouth University, Poole BH12 5BB, United Kingdom

^βBritish Geological Survey, Nottingham NG12 5GG, United Kingdom

Abstract

The emergence of the paradigm of the Internet of Things has underpinned the development of data-driven cyber-physical systems that collect and process data that is dense both in space and time. The application areas of such data-driven IoT systems are numerous and their socio-economic impact of great importance as they enable the monitoring and management of processes in sectors ranging from urban management to management of the natural environment. In this work, we introduce and detail an end-to-end technological framework for data-driven IoT systems for landslide monitoring. The framework is articulated in three tiers – namely *data acquisition*, *data curation* and *data presentation*. For each tier we present and detail its design and development aspects; from the IoT hardware design and the wireless communication technologies of choice, to how Big Data infrastructure and Machine Learning components can be combined to support a sophisticated presentation tier that delivers the true added value of a system to its final users. The framework is validated, extended and fine-tuned by means of two pilots at locations experiencing the impact of different landslide types and activity. This work qualitatively improves upon existing methods of landslide monitoring and showcases how data-driven IoT systems can pave new pathways for interdisciplinary research as well as generate positive impact on modern societies.

Keywords: Internet of Things; landslides; smart sensors; system architecture; LPWAN.

1. Introduction

Internet of Things (IoT) is an emerging technological paradigm that enables the massive and seamless interconnection of smart objects, things and machines over the Internet. The paradigm refers to an ecosystem of technologies applied in multi-sectors including geoscience, agriculture, farming, urban development, emergency response, medical, retail, security and surveillance that require continuous real-time data acquisition, analysis and evaluation.

The IoT paradigm is a key enabler towards the cyber-physical convergence. Small autonomous smart sensors and actuators deployed *en masse* enable the collection of *dense data both in space and time*. When coupled with technologies for Big Data analytics and Machine Learning, it allows the development of *data-driven cyber-physical systems* that can have a profound positive impact on individuals, societies and industries.

Thus far, the development of IoT systems has been primarily focused on the verticals of Smart Cities & Communities; Smart Manufacturing and Supply Chains; and Smart Agriculture & Farming. However, IoT systems can significantly contribute to the efficient monitoring and management of our natural environment; an area whose importance is highlighted ever stronger with the ongoing Climate Crisis [1]. Natural phenomena such as landslides can be destructive and can have profound negative socio-economic impact.

A landslide is defined as a movement of rock, earth or debris down a slope [2]–[4]. It is a physical system that develops through time in several stages driven by different phenomena including precipitation, snowmelt, temperature changes, earthquakes [5] and, if at the coast, marine processes. In the UK, the most active landslides are at the coast, driven by both terrestrial and marine processes. In other areas, such as India, precipitation levels are an important triggering factor for landslides.




Figure 1 Aerial view of the landslide site at East Cliff, Bournemouth, UK. Note the damages inflicted on the public access pathway (top of the cliff), the public access building (bottom of the cliff) and the historic Edwardian funicular (side of the landslide).

Climate change models forecast scenarios that indicate an increase in the triggering conditions for landslide occurrence (e.g. high-intensity precipitation events and sea level rise) that are indisputably linked to the stability of natural and engineered slopes [6].

As well as loss of life and physical damage, landslides can pose significant repeated socio-economic threats in landslide prone areas (e.g.[7]) where the built environment and infrastructure are particularly at risk in the long term. It is the potential for increased impact that has motivated the need for novel and efficient methods of landslide monitoring. Traditional landslide monitoring methods and techniques are limited both in technical terms (quality and frequency of data) as well as in usability terms (high inferred costs, difficulty of deployment, restricted access to areas of interest).



Figure 2 View of the public pathway leading to the beach front on top of the East Cliff, where the landslide incident took place in 2016. Red line denotes the new edge of the cliff after the incident.

1.1. Setting the Scene

The two pilot areas are on the south coast of England: Bournemouth and Barton-on-Sea. Both sites have landslides of different scales where significant damage to public infrastructure is reported.

East Cliff landslide, Bournemouth, Dorset. On 24th April 2016, a landslide occurred on the 35 m high section of the East Cliff in Bournemouth, Dorset (Figure 1). The day before, cracks at the top of the cliff had been reported to Bournemouth, Christchurch and Poole Council (BCP) (Figure 2) who cordoned off the area, both at the cliff top and at the foot of the cliff. The landslide damaged an Edwardian funicular railway as well as a toilet block and an array of fences and benches from the top of the cliff [8]. This event incurred high costs to the local authority for damage control and lengthy full-scale clean-up operations that spanned months. BCP council installed coir matting to minimise erosion and has employed several monitoring methods. These include conducting 6-

monthly topographic surveys, visual and digital inspections (using LIDAR), and groundwater monitoring.



Figure 3 View of the new cliff-top landslide site at Barton on Sea, Hampshire, UK. This is an active landslide posing significant risks to the public and near-by infrastructure and buildings.

Barton on Sea landslide, Hampshire. Much of the frontage at Barton on Sea is subject to ongoing complex landslide activity in the cliffs with significant events reported regularly. The area behind the cliff is continuously affected and is mostly residential with some businesses and infrastructure (Figure 3 [9]). The cliffs are 1.5 km long and 30-35 m high and have been subject to coastal protection and cliff stabilisation works since the 1930s¹ that have cost New Forest District Council (NFDC) significant amount of money. NFDC actively monitor and survey the landslide² with the support of the British Geological Survey (BGS) in the near future.

The limitations in relation to characterising the aforementioned monitoring and control methods include – among others, sparse data turnaround and high inferred costs – provided the motivation for researching and developing an IoT-enabled system for efficient landslide monitoring. The solution ought to be highly scalable, easily deployable and should infer low maintenance costs. Additional functional requirements included (as elicited from the stakeholders):

- The system should be able to monitor the landslide independently of any vegetation or other infrastructure (e.g. coir nets) present.

1

<http://www.newforest.gov.uk/CHttpHandler.ashx?id=21516&p=0>

- The system should be able to provide high fidelity data with centimetre-scale accuracy.
- The system should be able to allow remote monitoring of the landslide.
- It would be desired for the data transmission to be wireless.
- It would be desired for the data transmission to be (near) real time.

In this context, the proposal for a solution based on a framework combining IoT technologies for sensing and wirelessly transmitting data, Big Data infrastructure and Machine Learning for data curation and a sophisticated web-based user interface was put forward. Our methodological approach was first validated in a proof-of-concept pilot (referred to as Pilot 1) deployed at Bournemouth's East Cliff site that ran successfully for 6 months. An improved and significantly extended version of our system has been deployed in a second pilot (referred to as Pilot 2) at Barton-on-sea at a location that is closely surveyed by contemporary geological methods. The aim of Pilot 2 is to evaluate and compare the performance of the framework against these well-established monitoring methods.

1.2. Our Contribution

In this work, we present a technological framework for landslide monitoring which employs emerging technologies such as Internet of Things, Big Data analytics and Machine Learning. The framework has matured through and been validated in two Pilots; one at East Cliff, Bournemouth and one at Barton-on-sea.



Figure 4 The architectural tiers of the presented framework

The framework specifies a modular, loosely coupled architecture consisting of three tiers; *data acquisition*, *data curation* and *data presentation* (Figure 4).

2 <http://www.newforest.gov.uk/article/13525/Introduction>

Regarding the *data acquisition* tier, we detail the development of IoT sensing devices for landslide monitoring. We discuss the IoT boards that were used (a bespoke board for Pilot 1 and a commercially available one for Pilot 2), the wireless communication technologies used (SigFox for Pilot 1 and LoRa for Pilot 2) as well as the sensors used and their interfacing with the boards having in mind energy efficiency (detailed schemas are provided).

Regarding the *data curation* tier, we present the technologies used for collecting and handling IoT data streams. The backend architecture has been developed in a modular way having scalability in mind. Specialised technologies are used for Big Data curation and analytics as well as Machine Learning frameworks. The aim is for the backend not only to be able to handle streams of geological data from multiple sites, but also to be able to produce models that provide insights into landslide mechanisms as well as reliable predictions for anticipating future incidents.

Regarding the *data presentation* tier, the architecture and design of a web portal is presented. Its development has been based on functional requirements provided by specialists working on landslides (geotechnical engineers, council assistants and geologists). This front-end dynamically pulls data from the data curation tier and supports multi-modal data presentation featuring data overlaying with interactive maps and user-defined data plots.

The introduced data-driven IoT framework contributes to the following: factors

- It fills the apparent gap for a geological monitoring system that continuously provides reliable near-real-time data, and is independent of workforce restrictions, accessibility, weather conditions and limitations.
- It dramatically improves currently employed monitoring methods in terms of ease of deployment, scalability, cost efficiency, data precision, granularity of data, and agility of use for the end-user.
- It provides a state-of-the-art platform for real time data acquisition, processing and visualisation that benefits the geologists, local government, community and other stakeholders.

Finally, we would like to note that the framework does not intend to substitute the already existing methods, but aims at complementing them and, therefore augments the

arsenal of available methods. In particular, existing monitoring methods rely upon highly sophisticated and expensive equipment of high accuracy and precision that require (in several cases) the employment of corresponding specialised expertise. These methods, in principle, provide data of high fidelity that are sparse in space and time. On the contrary, the proposed framework makes use of highly affordable infrastructure that introduces low capital and operational costs, and that can be easily deployed. While individual devices are not characterised by the same levels of accuracy and precision, the system compensates this by making use of state-of-the-art technologies in Big Data analytics and Machine Learning.

2. Related work

2.1. Background on landslide monitoring.

Although, the primary driving force of landslides is gravity, the stability of a slope is influenced by a variety of other factors including: the material type and strength, lithological structure, hydrogeology, the slope angle, earthquakes, meteorological and other environmental conditions such as marine processes. Knowledge of these conditions can help to predict the location, types, and volumes of potential failures [10], [11]. The geological data related to landslides are mostly complex, and data acquisition-devices and real-time data transmission network infrastructures must be carefully designed.

Currently, available landslide monitoring techniques can coarsely be taxonomized in methods requiring physical access and methods of remote monitoring [12]. Topographic surveys and on-site visual inspection by experts are probably the most commonly employed method; however, this method is subject to ease of access and does not scale well with respect to area coverage. More accurate methods include the use of sophisticated equipment, such as terrestrial laser scanners [13], extensometers [14], tensiometers and inclinometers [15] (e.g. [16], [17]). However, these methods require specific expertise and introduce high operating costs. Furthermore, these methods are also not easily scaled to cover large areas. Remote monitoring methods include the use of satellites, either in the form of GPS systems [18] or interferometry techniques via synthetic aperture radar (SAR) satellite images [19]. It is worth mentioning that apart from the issues already discussed, few of the aforementioned methods are able to provide data in real

or near-real time. Also, some of these methods are not accessible to authorities of local communities either due to them requiring high expertise or due to inferred costs.

Landslides have been monitored throughout the world using a range of technologies over time (e.g. [16], [20], [21], [22]). As well as the aforementioned remote sensing techniques and visual inspections, traditional instrumented landslide monitoring systems collect information on landslide movement as well as conditioning factors such as weather conditions, groundwater, the geology and its geotechnical properties and, if at the coast, may include the state of the beach and any defences. Instruments include down-borehole inclinometers and piezometers, geophysics such as Proactive Infrastructure Monitoring and Evaluation (PRIME) systems [23] and weather stations [16]. Some of these technologies use telemetric systems but these can be expensive to install and maintain and may require connection to mains electricity for continuous use; many rely on manually downloading data on site and most require specific expertise.

2.2. Landslide monitoring using Wireless Sensor Networks.

Wireless Sensor Networks (WSNs) are peer-to-peer ad-hoc networks consisting of small autonomous sensing devices (a.k.a. sensor motes) that are able to carry out complex tasks collaboratively. WSNs are a key enabling technology of IoT [24], and their paradigm has contributed a lot in developing core IoT technologies such as IEEE802.15.4 and 6LoWPAN. There is abundant literature on WSNs being employed in several applications such as smart buildings [25], forest fire detection [26] and smart grids [27].

Regarding landslide monitoring with WSNs, the line of research presented in [28], [29] is probably the most notable one. The authors designed a column that houses several sensors for detecting landslides. In particular, the sensor suite consists of:

- *Dielectric Moisture Sensors*: Measure water content in the soil.
- *Pore Pressure Piezometers*: Measure groundwater pore pressure.
- *Strain gauges*: Measure movement of soil layers attached to the Deep Earth Probe (DEP).
- *Tilt-meter*: Measure the movement of soil layers regarding creep, slow or sudden movements.

- *Geophones*: Measure vibrations caused during a landslide.
- *Rain Gauges*: Measure the effect of rainfall on a slope and therefore, the ancillary effects such as pore pressure.
- *Temperature Sensors*: Physical properties of soil and water change with temperature, recorded every fifteen minutes.

The success of these devices was demonstrated by the early detection of a landslide in July 2009, providing validation for the authors' design during a substantial rainfall period in India's monsoon season. Criticisms of this system rely on the large physical form factor of the sensor columns (20 meters in length), their high energy consumption (the column relies upon a constant, wired power source which is backed by a power bank and a solar panel) and their high cost.

Following the emergence of Internet of Things as a novel ICT paradigm, monitoring systems for the physical environment started to employ relevant technologies and approaches. One such comparable IoT themed system exists in Panzhihua Airport, located in the Sichuan province, China [30]. Similarly to the framework presented in this work, this system's technical architecture also adopted a 3-layer approach; a sensor layer, a network layer and an application layer. The sensor layer of the system obtained and reported information about the physical geography of the slope, such as landslide deformation information like crack displacement, GPS and angle deformation and local rainfall data. The network layer included private networks employing cabled and wireless technologies, a network management system and an Internet-based cloud computing platform. The later was responsible for the transfer and processing of data from the sensor layer. The system utilised a total of 5 extensometers deployed in strategic placements along a fault crack in the slope, with one end on the landslide and the other on firmer ground.

While the system successfully issued warnings for an impending landslide, it demonstrates several limitations compared to the framework presented in this work. Firstly, it comprises specialised equipment (such as extensometers) that require specialised expertise to be deployed; they also increase the overall cost of the system significantly. In our approach, we employ consumer-graded electronics that are abundantly available in the market, thus greatly reducing incurred costs and facilitating development and deployment. Secondly, we employ LPWAN communication technologies that allow

for seamless device interconnection to the Internet, therefore reducing the complexity of the system, increasing the battery lifetime of the devices, and further reducing incurred costs. Finally, the presented framework provisions the use of Big Data Analytics and Machine Learning infrastructure in the back-end, therefore greatly increasing the scalability of the system, enabling the system to support multiple sensing modalities and providing deeper insights to the geological processes that trigger landslide incidents.

In [31] authors present a landslide monitoring system that makes use of wireless sensor networks and the LoRa LPWAN. The focus of this work is on the precise geolocation of the nodes of the network, such that the motion of the slope is monitored in detail. To this end, authors provide details on the sensors they employ (accelerometers, magnetometers, angular rate sensor) and the network architecture, consisting of three anchor nodes located at the top of the slope that enable the triangulation of the position of each node. While this system also uses LoRa and low-powered single board computers, in contrast to the presented framework in this work, it does not provide for Internet connectivity and does not support any sophisticated infrastructure for curating and leveraging upon the collected data. Finally, authors provide little insights on the set up of the sensing devices and their expected efficiency in energy consumption.

In [32] authors present an IoT-based monitoring system employing two different node types; a meteorological node and a ground monitoring node. Similarly to our approach, the system makes use of consumer-graded electronics, such as LoPy [33] – a developing IoT platform supporting multiple wireless communication technologies. The system is detailed in terms of its back-end infrastructure, which comprises a lambda architecture that communicates over a web interface with a scientific sharing platform. Although not explicitly mentioned, the architecture is capable of supporting the training of Machine Learning models for landslide forecasting. However, contrary to our presented framework, in [32] authors do not include explicitly in their design a front-end component. This is a commonly overlooked, yet crucial, component since it allows the end-user and the system stakeholders to harvest the added value of the system and collected data, e.g. by means of Visual Analytics. Furthermore, in [32] authors do not elaborate on the hardware configuration of their sensing devices and their energy efficiency. This is a crucial aspect of a monitoring system for the natural

environment, since deploying and reclaiming devices can be very costly and challenging.

2.3. Background on IoT technologies.

Technological advances have greatly promoted the IoT paradigm in recent years. Regarding IoT hardware, the emergence of Single Board Computers, such as the




Figure 5 Comparison of Sigfox, LoRa and NB-IoT[43]

Application Layer					3GPP
CoAP, AMQP, MQTT, MQTT-SN, XMPP, DDS, WebSocket...					
Service Discovery: mDNS, DNS-SD, SSDP, SLP...					
Security	TLS, DTLS				
Transport Layer					
TCP / UDP					
Internet Layer					
Addressing: IPv4 / IPv6		Routing: RPL, CO RPL, CARP...			
Adaptation		6Lo WPAN, 6TISCH, 6Lo, IEEE 1095.1...			
Network Interface					
IEEE 802.15.4 (LR-WPAN)	IEEE 802.15.1 (Bluetooth)	LPWAN (LoRa WAN...)	IEEE 802.15.6 (WBAN)	RFID / NFC	
IEEE 802.11 (Wi-Fi)	IEEE 802.3 (Ethernet)	IEEE 802.16 (WiMax)	Z-Wave	IEEE 1901.2 (PLC)	

Figure 6 The ecosystem of IoT technologies [44]

Table 1 List of terms and acronyms used

Key Term	Definition
6LoWPAN	IPV6 over Low Power Wide Area Network. Allows the sending of data using IPV6.
Adafruit	An electronics company specializing in single board computers.
Amazon Web Service (AWS)	An Amazon solution providing on demand cloud services and API's to users over the internet.
Application Programming Interface (API)	An interface or communication protocol between different parts of a computer program or system.
BlueFox v2.7	A single board computer provided by NetSensors Ltd.
Bootstrap	Framework for the quick development of front-end web application user interfaces.
Django	A python-based web-framework, allowing for the hastening of web application-based programming.
IEEE802.15.4	Technical standard for low-rate Wireless Personal Area Network.
Java Script	High-level, just-in-time, programming language.
JSON (JavaScript Object Notation)	A file format allowing the communication of data objects.
Kafka	Used in the development of real-time data pipelines.
LoRa (Long Range)	A low power wide area network technology for IoT, with an option for personal network ownership, or to join a greater network.
LoRaWAN	LoRA Wide Area Network. A network created by the intercommunicating of LoRa devices.
LoWPAN	Low Power Personal Area Network.
LPWAN	Low Powered Wide Area Network.
MCU	Micro-controller unit.
Model, View, Controller (MVC)	Style of web application structure. Splits a web application into 3 parts, models, views and controllers.
NB-IoT (Narrowband-IoT)	Low Power Wide Area Network designed for IoT, offering improvements to power efficiency.
Node-RED	Flow based visual programming software.
opAmp	Operational Amplifier.
Python	4 th Generation programming language.
SigFox	A narrow band, low power wide area network technology for IoT with paid access.
Sparkfun	An electronics company specializing in single board computers.
Zigbee	Low Power Personal Area Network, using IEEE802.15.4 Standard.
Z-Wave	Wireless protocol, mainly used in home automation with a range of 100m.

Raspberry Pi and the Arduino, have popularised IoT enabling the development of “at home” projects by non-specialists. This has led to the growth of the corresponding market with new vendors, such as Adafruit and Sparkfun, introducing novel board designs while reducing the costs of purchase.

Corresponding advancements have taken place regarding the wireless communication technologies. The landscape of IoT wireless technologies initially was dominated by Low Power Personal Area Networks (LoWPANs) making use of IEEE802.15.4, such as Zigbee, Z-Wave

and 6LoWPAN. This family of protocols was the first to address the device characteristics imposed by the IoT paradigm directly (e.g. their highly constrained nature in computational resources and connectivity), however, they supported communications over small distances (in the order of 100 metres). This necessitated the use of gateways, which in turn hindered the scalability of IoT systems.

The introduction of Low Power Wide Area Networks (LPWANs) substantially changed the landscape. By operating at low sub-GHz frequencies, LPWANs support

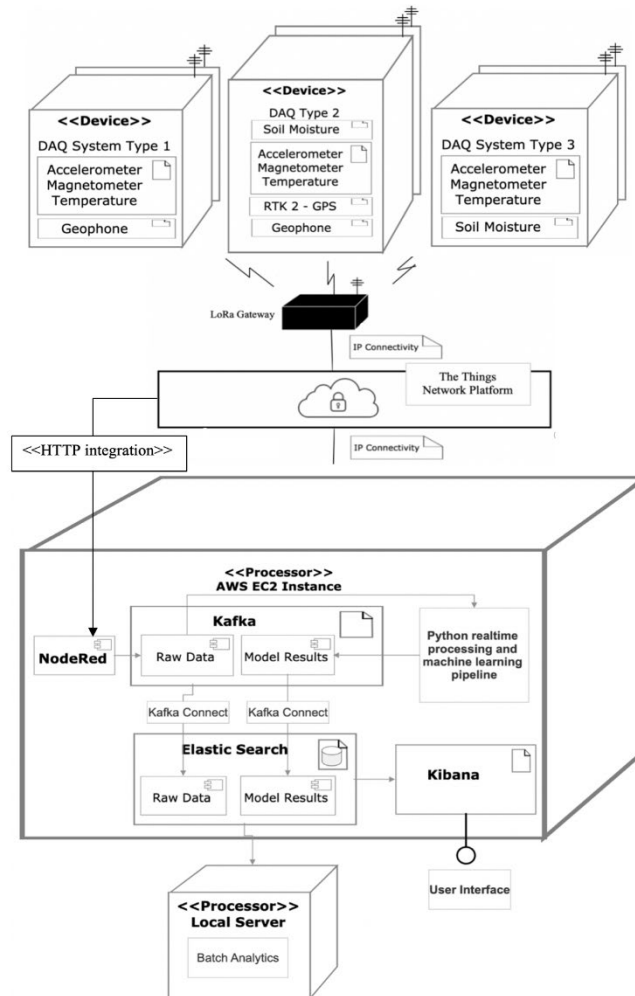


Figure 7 The reference architecture of the framework for landslide monitoring

long range wireless communication (in the order of tens of kilometres) while still being highly energy efficient. In spite of this being achieved at the cost of low baud rates, they achieved network speeds suffice for most IoT applications. Furthermore, LPWANs successfully mitigate the issues of scalability, thus paving the way for the deployment of large-scale IoT systems. Figure 6 illustrates the layers of IoT protocols prescribed by IEEE and ETSI (European Telecommunications Standards Institute).

Sigfox [16] was among the few first LPWANs to be commercially available. The network is organised in cells, each one covering a maximum area of 50km, allowing up to 140 uplinks and 4 downlinks per device,

per day. The technology operates as a one-hop star topology although the network requires a mobile operator to carry the generated traffic. In the Sigfox business model, the network access points are owned by the Sigfox company or official representatives, and users need to pay a premium per device per day.

LoRa [17] – abbreviation for Long Range – is another popular LPWAN. Similarly, to Sigfox, it operates at low, sub-GHz frequencies at low duty-cycles, thus addressing the range-vs-energy trade-off. The main difference to Sigfox lies in the business model it assumes. The ownership and operation of LoRaWAN access points are open to everyone, and therefore, LoRa supports the operation of private IoT networks. Certain initiatives,

such as The Things Network, leverage upon this model in order to develop crowdsourced IoT networks spanning across entire regions (the UK network spans across the entire country and continues to grow).

Complimentary to Sigfox and LoRa – wireless technologies that rely upon IoT-specific infrastructure – there are IoT wireless technologies being introduced that operate over the existing cellular network. NB-IoT (narrowband IoT) is such a technology that focuses specifically on indoor coverage, low cost, long battery life, and high connection density. It makes use of the LTE standard but dramatically limits the bandwidth, thus achieving energy efficiency. The great advantage of this technology is the fact that it makes use of the existing cellular infrastructure that already provides good coverage both indoors and outdoors. Figure 5 demonstrates a comparison of Sigfox, LoRa and NB-IoT. Table 1 summarises the terms and acronyms definitions used.

3. Overview of the Landslide Monitoring Framework and of the Two Pilots.

Figure 7 depicts the reference architecture of the presented framework. It is articulated in three tiers that are loosely coupled with the use of APIs. This provides flexibility in fine tuning and adjusting the framework by amending the design of each tier and the choice of identical technologies used independently of the rest of the system. The *data acquisition* tier consists of the IoT sensing devices that are fitted with specialised sensors and are deployed in an area of interest. The devices are autonomous and highly energy efficient, being able to continuously operate for an extended period of more than 600 days (see detailed discussion in Section 4.2). The devices collect geological data, which is then transmitted over a LPWAN network to the next tier. The *data curation* tier employs technologies for efficiently storing and managing IoT Big Data streams, as well as for training Machine Learning models for landslide monitoring. The architecture assumed is highly scalable and can accommodate (and if needed aggregate) data streams from multiple sites as well as from external sources, such as weather data. Finally, the data presentation tier dynamically retrieves data from the data curation tier visualising them using several modalities, such as overlaying data on Google maps, dynamic selection of data categories and figure plotting by the end user, cross-referencing data from multiple sources, and so

on. Each tier is presented in detail in the following sections.

Design and development for Pilot 1 commenced after a landslide incident that took place in 2016 at Bournemouth's East Cliff. The morphology of the




Figure 8 One of the sensing motes used in Pilot 1 in its casing. The device features among other, a dual microprocessor, SigFox connectivity, a 3-axis accelerometer, soil humidity and temperature sensors and is powered by two 4200mAh batteries (visible in the back of the device).




Figure 9 The sensing device in its final form prior to being deployed at the landslide site. The soil humidity and temperature sensors are visible, waterproofed via silicon glands. Installation was undertaken by specialist rope technicians and geotechnical engineers.



Figure 10 Testing connectivity of Pilot 1 device prototype. Initially, connectivity was tested with the local SigFox AP at the base of the East Cliff landslide at Bournemouth beach (sub-fig.(a)); then connectivity was tested at a depth of 10 inches to test signal soil penetration (sub-fig. (b)). Sand was then saturated with water to confirm signal penetration through wet sand at the same depth (sub-fig. (c)). Note that wet sand is one of the most harsh media for wireless signal propagation. Finally, verification of packet reception on the dashboard of the Sigfox network server (sub-fig (d)).



Figure 11 Aerial view of the landslide site East Cliff, Bournemouth, UK. Red symbols denote the deployment locations of the devices (symbols are also used to identify each device). One device (denoted by circle) is deployed on a mild slope; another one (denoted by cross) on a steeper slope; while the third one (denoted by dash) is deployed on the stone-paved support structure.

surrounding area in Bournemouth is characterised by cliffs that overlook a ten-mile stretch of public-access beachfront. The type of soil and local weather conditions – strong winds, high air humidity, precipitation levels during winter and high temperature variations during summer – make the area prone to landslides. Furthermore, landslide incidents pose a significant risk to public safety, as the cliffs and seafront are located well within the urban fabric of the town. The aim of Pilot 1, on one hand, has been to showcase the capabilities of data-driven IoT-enabled systems to the relevant stakeholders and, at the same time, validate the use of relevant technologies in application areas that move well-

beyond the usual domains – Smart Cities, Smart Manufacturing, Smart Agriculture – and into the domain of environmental sciences. In this context, validation entailed not only the definition of a data-driven system architecture for IoT systems, but also significant engineering challenges related to the choice of sensors, energy efficiency and the physical design of the sensing devices so as to withstand weather conditions and other adversarial elements (e.g. seagulls and rodents are known hazards for any equipment deployed on sea-side cliffs). To this end, three devices were deployed on site at carefully selected locations, each one demonstrating different characteristics (see Fig. 11) – one on a mild




Figure 12 The landslide at Barton on Sea, Hampshire, UK, showing the approximate deployment location of the sensing devices.

slope; another one on a steeper slope; while a third one on the stone-paved support structure. This enabled us to calibrate and validate the configuration of the devices by monitoring locations with different expected geological behaviour.

Following the successful conclusion of the first pilot, the aims of Pilot 2 were firstly, to leverage upon lessons learnt from Pilot 1 and improve the framework accordingly and, secondly, compare the quality and value of data provided by the system to data collected via well-established and proven methods employed in Geology. For these reasons, the advice of expert geologists of the BGS was sought, and the location of Barton on Sea was selected. At Barton on Sea there exists an active landslide that due to its on-going impact on the local community (public infrastructure has already been affected and private property may be at risk in the future) is actively monitored by local authorities and the BGS. This makes the location ideal for trialling and further evaluating the framework. A reconnaissance field survey was carried out in July 2019 and deployment positions were agreed to place the sensors. As the site is publicly accessible (the base of the landslide is part of a public beach), the exact locations must remain undisclosed for security reasons. The exact locations must remain undisclosed for security reasons, but they are chosen to maximise the likelihood of capturing landslide movement, facilitate installation

and be at a location monitored regularly by the authors and collaborators throughout the project.

4. The Data Acquisition Tier

The main functionality of this tier is to collect and transmit sensory data from the area of interest to the data curation tier. The tier includes the sensing IoT devices (sensors and IoT boards) as well as the IoT network. Regarding the devices, design considerations include the minimum required technical specifications of the boards (these are usually considered in the context of a trade-off with energy efficiency), means of powering the devices, the suite of sensors the devices fitted with and, finally, the physical form factor. The choice of wireless communication technology is dictated, on the one hand, by the functional requirements of the system (mainly volume and rate of data) and on the other hand by availability of and access to relevant infrastructure as well as on inferred capital and operational costs.

4.1. Data acquisition in Pilot 1

The IoT boards developed for Pilot 1 were based on the BlueFox v2.7 platform, provided by Net Sensors Ltd (a Bournemouth-based SME). The initial rationale was to use a bespoke hardware platform to maximise energy efficiency. The boards featured two low-power microprocessors, a SigFox modem, a 3-axis digital output accelerometer, humidity and temperature sensors. The




Figure 13 Wiring diagram of the IoT boards for Pilot 2 equipped with GPS-RTK and geophone. The wiring permits the geophone to continuously monitor seismic activities and in the case of an incident to generate a wake-up signal to the board. This allows the board to remain in deep sleep mode for energy efficiency.

sensing capabilities of the boards were extended by means of a DS18B20 Waterproof Digital Thermometer [34] and the Analog Capacitive Soil Moisture Sensor V1.2 by DFRobot [35]. Finally, each device was powered by two D-cell batteries at 4200mAh capacity each, giving a total capacity of 8400mAh (Figures 8 and 9). The devices were programmed in C using the Arduino compiler/IDE.

The choice of Sigfox as the LPWAN technology to be used was due to its availability in the area. In particular, Bournemouth is one of the focus areas in the UK for deploying and promoting the use of IoT technologies under the Things Connected³ initiative of Digital

Catapult UK, and Sigfox provide a 100% coverage in the area. Although the use of the Sigfox network requires paying a premium per connection, this cost was covered by Digital Catapult.

Since the motes were to be deployed in an outdoor environment, exposed to weather conditions and other hazards (seagulls and rodents have been proven to pose great threats for any type of equipment), particular care was taken to protect the motes. For this reason, the boards were waterproofed by the application of an acrylic conformal coating and an epoxy resin to protect the circuit boards of the capacitive sensor, and they were encased in an IP-67 graded case.

³ <https://www.digicatapult.org.uk/projects/things-connected/>

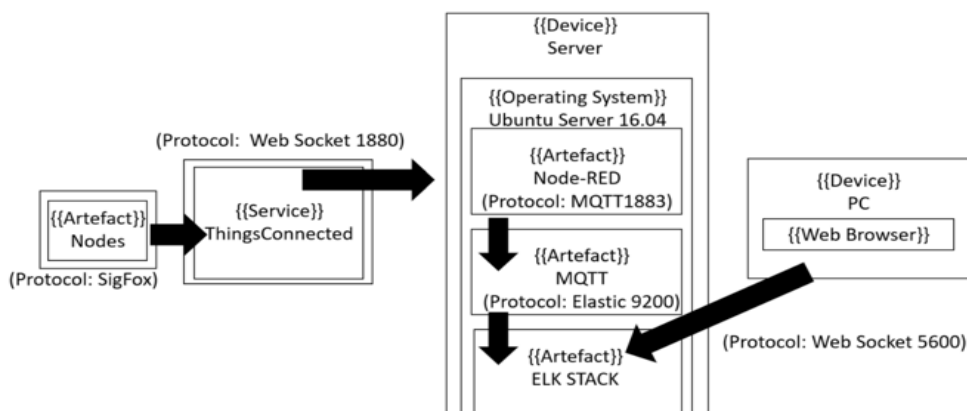


Figure 14 The Architecture of the Data Curation for Pilot 1.

Prior to the final deployment, extended tests were carried out in order to investigate the penetration rate of the Sigfox signal to the ground and to establish the ability of the devices to successfully transmit sensory data. The first consideration had to do with the morphology of the area of interest and the possibility that the overlooking cliffs would cast a “connectivity shadow” over the devices. The second consideration had to do with the ability of the devices to wirelessly transmit through wet sand (Figure 10) - it is worth noting that this question is not well studied in the literature. The tests were successful, and the devices were deployed by cliff-hanging specialists and operated continuously for 6 months while collecting and transmitting data at 15-minute intervals.

4.2. Data acquisition in Pilot 2

Building upon lessons learnt from Pilot 1, the design of the sensing IoT devices for Pilot 2 was revisited. For the second pilot an off-the-self board was selected (Sparkfun Pro SAMD21 [36], [37]) in order to benefit from the wide development support from the manufacturer and the community. The sensor suite was also revised after consultation with experts of BGS. The new revised suite consisted of the following components:

- A 6-Degrees-of-Freedom (6DOF) inertial measurement unit (IMU) including an accelerometer, a magnetometer and a thermometer (board LSM303C [38]);
- A high-fidelity geophone providing seismic data in 1D, 2D and 3D seismic surveys (model SM-24 [39]);

- An analogue capacitive soil moisture sensor (model DFRobot SEN0193 [35]);
- An RTK2 (Real Time Kinematics) GPS board from ublox [40] for high-precision location data (accuracy ~1cm) for monitoring soil movement.

The sensors were interfaced with the IoT board using the I2C protocol and the QWIIC cable interface. The geophone made use of its own breaking board and two operational amplifiers. This configuration enabled the IoT board to remain in deep sleep state (thus greatly reducing energy consumption), while the geophone was continuously monitored for seismic activity – hardware frequency filters made sure the geophone was monitoring only on the desired spectrum of seismic activity, thus eliminating any false positives. Should any activity be detected, the geophone triggers an interrupt that would wake up the device. The wake-up time – consisting in the geophone picking up vibrations at a certain frequency spectrum and the break-out board generating an interrupt to the main board - is in the order of 30ms; quick enough to capture any ongoing incidents. Details of the wiring of the devices (version with geophone and the RTK-GPS module) are depicted in figure 11.

Regarding the LPWAN technology, the location for Pilot 2 did not provide Sigfox coverage. Furthermore, the restrictions imposed by Sigfox both on the maximum number of transmissions per day (144 messages) and the maximum size of each transmission were deemed rather restrictive for the application of landslide monitoring; particularly in the case of an incident when sensing and data transmission would need to be very dense. Therefore, for Pilot 2, the technology of choice has been

the LoRaWAN. LoRaWAN also carries several advantages. Firstly, one can deploy and manage their own private LPWAN network by procuring the required infrastructure (i.e. the wireless access points) privately. This is of great importance for the presented framework, as it enables authorities, such as local councils or the BGS, to commission their own networks. Furthermore, this greatly reduces the operating costs as LoRaWAN, contrary to Sigfox, does not require a subscription. Finally, LoRaWAN initiatives, such as The Things Network, support the adoption of crowdsourcing practices allowing the general public or third parties to augment the network's coverage by providing roaming services via their own access points.

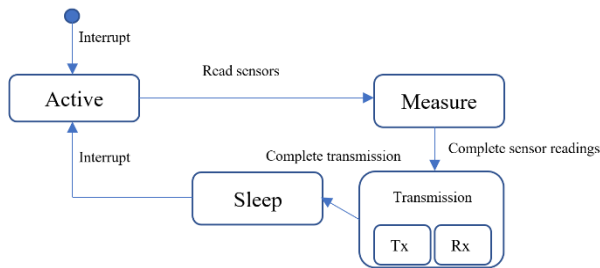


Figure 15 State transition diagram capturing the operation states of the sensing devices.

Regarding the energy source, each device is equipped with two D-cell Li-Ion 3.3V batteries at 4200mAh each, totalling 8400mAh. Regarding battery lifetime special care has been taken to maximise energy efficiency and reduce energy consumption. This entailed both hardware interventions (such as dismounting LEDs off the main board and using programmable hardware switches to drive sensor components) and software optimisations. Figure 15 depicts the states and the related events of the activity cycle of the sensing devices. The states include the device waking-up, being active, conducting sensor measurement, transmitting data (comprising both transmit and receive modes), and finally the sleep state. Following is an operation outline for each state (time interval and drawn current figures were taken by means of specialised multi-meter).

Active state – This state entails powering up and operating the device components that are needed to support basic functionality of the device; i.e. the MCU, the LoRa RF module, the operational amplifier (opamp) breakout board, etc. At this state, other sensors (such as the accelerometer and the soil moisture sensor) are kept on the power-down state by means of hardware switches to increase energy efficiency.

Measure state – This state entails also powering up all sensor components, reading and registering sensor measurements. These include the 6-axis digital accelerometer, the soil moisture sensor, the opAmp breakout board and the geophone. The geophone itself is a passive unit (i.e. it does not draw any current), however, the opamp breakout board that drives it consumes a considerable amount of current. To mitigate any spikes in power consumption, sensors are sequentially powered up by means of a programmable hardware switch and powered down once the sensory readings are registered.

Transmission state – This state entails the transmission of the registered sensory readings over the LPWAN. It consists of both transmit and receive modes for the radio module, also accounting for the reception of acknowledgments. The current drawn in this state can vary based on payload size, data rate, and selected spreading factor (for the LoRaWAN network). The power consumption can be reduced by making the right choice of gateway location, payload design and deployment model. In this state peripherals and sensors are powered down to further increase energy efficiency.

Sleep state – In this state the entire device enters an extremely low power consumption mode where all components (including the MCU and radio module) are powered down. The device remains in this state the grand majority of the time thus achieving an extremely low duty cycle. The device exits this state when a corresponding interrupt is generated (e.g. when a timer expires).

In order to estimate the expected battery lifetime of the devices, we have carried out detailed measurements using specialized equipment in our electronics lab and an analysis of the power consumption of the devices. In particular, by considering the time duration and current consumption at each state, we calculated the total charge (measured in mAh) consumed during each state per operational cycle. Having configured the devices to collect sensory measurements periodically every 3 hours, we calculated the estimated number of cycles and, therefore, the time each device will operate before depleting its batteries to be slightly more than 600 days. Table 2 summarises our analysis.

5. The Data Curation Tier

Once one of the IoT sensing devices collect some sensory data, this data is then transmitted over the LPWAN network to the corresponding network server and from there to the data curation tier. This tier of the framework

Table 2 Summary of measurements and analysis regarding the expected battery lifetime of the sensing devices.

C : Total available battery charge (mAh)	C = 8400mAh
Active State	
T _a : Amount of time the device remains in active state (ms)	T _a = 0.1ms
I _a : Active state current consumption (mA)	I _a = 9.5mA
Z _a : Charge consumed in active state (mAh)	Z _a = 0.03x10 ⁻⁶ mAh
Measure State	
T _{geo} : Amount of time for taking and registering geophone readings (ms)	T _{geo} = 20ms
I _{geo} : Current consumption during geophone readings (mA)	I _{geo} = 9.5mA
Z _{geo} : Charge consumed during geophone readings (mAh)	Z _{geo} = 52.78x10 ⁻⁶ mAh
T _{sen} : Amount of time for taking and registering accelerometer and soil moisture sensor readings (ms)	T _{sen} = 1ms
I _{sen} : Current consumption during accelerometer and soil moisture sensor readings (mA)	I _{sen} = 13.82mA
Z _{sen} : Charge consumed during accelerometer and soil moisture sensor readings(mAh)	Z _{sen} = 3.84x10 ⁻⁶ mAh
Z _{meas} : Total charge consumed in measure state (mAh)	Z _{meas} = 56.62x10 ⁻⁶ mAh
Transmission State	
T _{tx} : Amount of time for transmitting data (ms)	T _{tx} = 73.3ms
I _{tx} : Current consumption during transmission (mA)	I _{tx} = 32.1mA
Z _{tx} : Charge consumed during transmission (mAh)	Z _{tx} = 654x10 ⁻⁶ mAh
T _{rx} : Amount of time for receiving data (ms)	T _{rx} = 1ms
I _{rx} : Current consumption during reception (mA)	I _{rx} = 15.8mA
Z _{rx} : Charge consumed during reception (mAh)	Z _{rx} = 4.38x10 ⁻⁶ mAh
Z _{trans} : Total charge consumed in measure state (mAh)	Z _{trans} = 658.38x10 ⁻⁶ mAh
Sleep State	
T _s : Amount of time the device remains in active state (hours)	3h
I _s : Current consumption during sleep state (mA)	0.57mA
Z _s : Charge consumed during sleep state (mAh)	Z _s = 1.71mAh
Total charge consumed per 3-hour cycle	
Z _{total} = Z _a + Z _{meas} + Z _{trans} + Z _s	Z _{total} = 1.710715mAh
Total number of cycles of operation; each cycle lasts 3 hours (C / Z _{total})	4910 cycles ≈ 14730 hrs ≈ 613 days

effectively amounts to a back-end Big Data storage and processing system leveraging Big Data Analytics and Machine learning technologies.

5.1. Data curation in Pilot 1

Pilot 1 saw the development of a simple back-end system, serving as a proof of concept for the landslide monitoring framework and giving space to experiment with technologies. Figure 12 depicts the corresponding reference architecture, consisting of a cloud-based server (an AWS m4.large instance running Ubuntu Server 16.04), hosting a Node-RED instance, an MQTT Server and Broker, and the Elastic stack.

The network server for both Pilots (Pilot 1 using Sigfox and Pilot 2 LoRaWAN) was provided by Things Connected; a Digital Catapult project initiative. Node-RED is a flow-based development tool developed by IBM for connecting hardware devices and APIs towards the IoT paradigm [12]. It is implemented in JavaScript utilising the Node.js framework, establishing a dataflow driven design tool that consists of JSON (JavaScript Object Notation) data generated internally and externally from the application. Node-RED was employed to provide a means of establishing a modular framework for future extensions and providing the means to export the data with ease towards the Elastic Stack.

Message Queuing Telemetry Transport (MQTT) is a lightweight publish and subscribe protocol for IoT and resource-constrained devices. MQTT was chosen in similar regards to providing a modular system design, by publishing topics, such as landslide data, and subscribing to these so that they may be logged and sorted within the Elastic Stack. The MQTT client is featured as a part of a Python Script, which specifies which topics to subscribe to, and converts the timestamps and JSON packet into a searchable variable within Kibana.

The Elastic Stack, or ELK stack, is the terminology to define three open source projects, namely Elasticsearch, Logstash and Kibana. Elasticsearch is a distributed, RESTFUL search and analytical engine as JSON over HTTP, and excels at indexing large amounts of text. Logstash manages events and logs, collecting them and parsing them for storage and layer usage with the two complementary technologies. Kibana is a data visualisation and exploration tool that acts as a dashboard for the stack and allows the creation of graphs and figures to express stored data.

Throughout the course of the first pilot, this first approach of the data curation tier proved to meet the functional requirements of the framework. The system was able to manage the ingress flow of IoT data streams. It provided quick search functionality and access to stored data and by leveraging upon Kibana, it also supported basic data visualisation (more sophisticated functions are provided by the Data Presentation tier).

5.2. Data Curation Pilot 2

The Data Curation tier for Pilot 2 built upon the first iteration, revamping the architecture and adding Big Data Analytics and Machine Learning technologies. This second iteration streamlined and improved the architecture put in place in the Pilot1, offering a more mature, modern and effective Data Curation tier.

Much like Pilot 1, the backend system is hosted on a cloud server (an AWS Elastic Compute Cloud (EC2) instance). The core technologies used in the first iteration were maintained (Things Connected network server, Node-RED and Elastic Stack). The data processing and analysis components that were introduced include Kafka and a Python-based analytics system. The software architecture is depicted on figure 17.

Raw sensory data arrive in the system via an HTTP interface from the 'The Things Network' server. Node-RED captures and processes the data payload, parsing it into data segments and performing an initial pre-processing before being fed to Kafka. Kafka is a real-time data processing and messaging platform used to send data to machine learning models periodically.

From Kafka, data are consumed by two processes: Elasticsearch and a Python script. Elasticsearch indexes the data, and the Python script begins the data analysis process. Python-based machine learning models are used on time-series data to accomplish predictive tasks. Python libraries such as Pandas, Scikit learn and TensorFlow are used here. From the Python Script, data is used to train a set of prediction models in real-time. The results of the model are sequentially published back into Kafka, another Elasticsearch instance consumes them.

Kafka is integrated into the system using the Confluent Platform. This is a complete streaming data platform, providing a full production-ready solution for real-time data streams, with Kafka at its core. This platform was chosen for its use of Kafka, easy integration and high-performance clients for Kafka API [41]. Elasticsearch acts as the primary data storage for the backend. Multiple

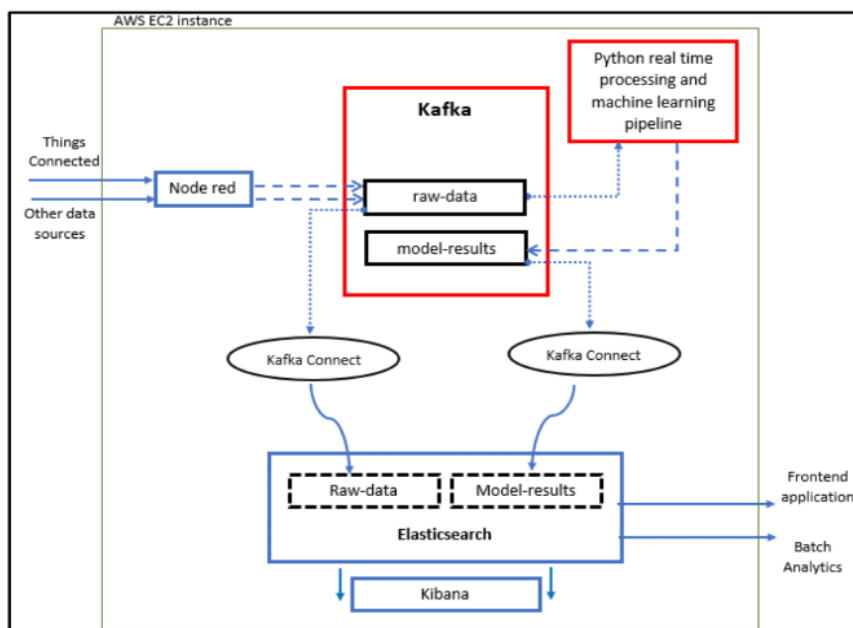


Figure 16 The Architecture of the Data Curation for Pilot 2

indexes are defined to hold different information such as sensor data, model results and logs. Kibana is used to manage the Elasticsearch cluster and visualise sensor and model data.

6. The Data Presentation Tier

An aspect that is commonly overlooked in the design and development of data-driven IoT systems is that of data presentation. Indeed, for systems whose primary function is the collection and use of data (such as monitoring systems) their true added value resides in the final end-users of the system being able to efficiently and conveniently consume collected data and information. How this will be achieved is a complex and multifaceted challenge that strongly depends on the nature of the application and the type of data, as well as on the stakeholders and end-users of the system. In the case of landslide monitoring the aim has been to develop a presentation layer that would be agile enough to accommodate the needs of both high-expertise users (e.g. geologists and geotechnical engineers) and local community stakeholders (e.g. local governments and planning authorities).

The data presentation tier utilises a front-end system that integrates with the back-end infrastructure to display collected and curated data to the end-user in a meaningful and impactful way. In particular, in the presented framework, it took the form of a Web Application created using the Django web framework and written in Python. The UI elements of the system were created using the help of Frameworks such as Bootstrap, enabling rapid development of a UI.

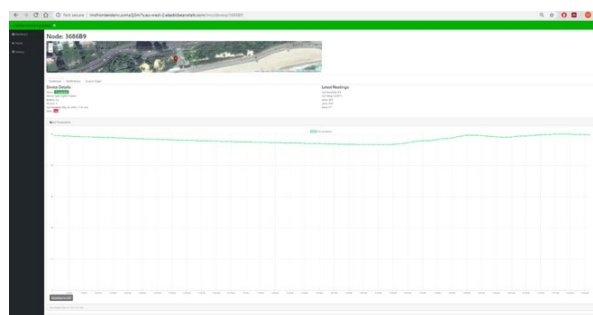


Figure 17 Data Visualization 'Devices' Page

Django is a widespread high-level python framework that uses a Model View Controller (MVC) design. It encourages rapid development by taking care of the hassle in web development. It handles the 'public'

features allowing a developer to purely focus on business logic. Further, Django has been designed to maximise speed of development, application security and scalability [42]. Data displayed and used on the front-end environment are data made available by the back-end system and data analytics pipeline. The web application links directly to the Elasticsearch instance present in the data curation tier. This is accomplished with the use of the Python library `Elasticsearch_DSL`. This is used to create Django style Get Method models that retrieve sensor data stored on the backend Elastic Stack just like any other Django Model to database interaction. `Elasticsearch_DSL` provides a more convenient and friendly way to write queries to Elasticsearch, it keeps close to the format of Elasticsearch querying, keeping its terminology and structure, whilst allowing its expression in the Python language.

This intuitive web application contains a full arsenal of data viewing and manipulation methods such as; Interactive Maps integration, enabled by the Python Library Leaflet JS; Custom graphing, using JS graphs; Kibana integration; Device specifications and statistics; and a notification system, with fully customisable user-defined thresholds. All data handled by the front-end system is real-time, allowing the web application to express a constant and continuous up-to-date state of landslide sites, thus allowing trends to be observed as they appear.

The maps integration allows the depiction of device nodes onto an accurate World map, allowing the presentation of physical device locations, in relevance to each other. This is coupled with the interactive ability to manipulate the maps scale and viewable region and functionality to select specific nodes present on the map. This also opens a page specific to that device, where specifications and statistics on the device can be seen, as well as custom and generic graphs explicitly relating to that device's data. Custom graphing is achieved by JavaScript code, and allows for the generation of graphs with selectable data sources and is the primary way the application displays data to an end-user. `Elasticsearch_DSL` queries populate the graphs with data stored on the Elasticsearch instance, this is updated in real time, providing the user with a constant, up-to-date stream of data.

In order to not lose functionality provided by the backend implementation, the web application is linked with a Kibana dashboard created as a product of the back-end Elastic Stack. This snapshot is shown to the user and

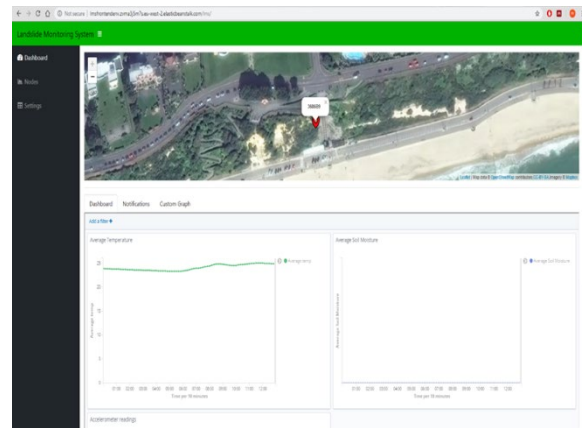


Figure 18 Data Visualisation ‘Dashboard’ page

allows them to keep functionality and visualisations created on Kibana, whilst hiding developer functions. Additionally, using `Elasticsearch_DSL` search queries, device specifications and statistics are printed to the user.

The above features provide an end-user with a plethora of manipulation and viewing methods, providing versatility and diversity in the application's use. Notification and warnings are also included; these send site-wide warnings when user-defined threshold levels have been met or exceeded for selected variables.

The functionalities supported by the data presentation tier are the result of a designing stage that included eliciting requirements and needs from stakeholders. Its four main sections are: Dashboard; Devices; Node; and Settings. It is worth noting that the sections have feature and functionality overlaps, but present their functionality in differing ways; for example, Dashboard shows high-level information, whereas Devices offers specific in-depth information.

6.1. The “Dashboard” section

This section serves as the main ‘Hub’ of the tier. It includes the interactive map that shows the location of all sensors, and these sensors themselves can be interacted with to take the user to a device's specific page. Additionally, using the applications notification functionality, if a node on the map has an alert (e.g. because soil activity was detected), then the node changes colour to red, visually notifying the user of an error or alert with that device.

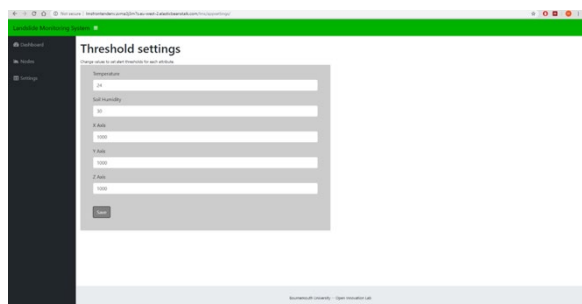


Figure 19 Data Visualisation ‘Settings’ Page

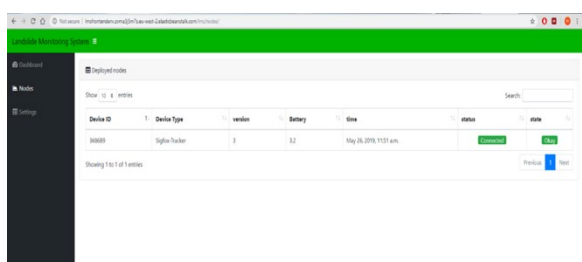


Figure 20 Data Visualisation ‘Nodes’ Page

Below this feature, further content is split into three tabs. The first tab shows all sensor data in graphs, per sensing device. The next tab contains all the notifications of the system, and these notifications are alerts that have been created due to incoming values exceeding set thresholds. The final tab includes a custom graph that allows the users to customize what data is shown and from what node.

6.2. The “Devices” section

This section includes details of specific devices, each device having its own dedicated view. Alongside this is a list of device specific details as well as a list of the latest figures are shown to end-users. Below this, there is a 3-tab content system: the first tab displays readings from the device in graph format; the second tab includes notifications specific to that device, and the final tab includes a custom graph that allows the user to manipulate what data is shown.

6.3. The “Settings” section

This section includes settings that allow the user to adjust the alert threshold levels based on sensor readings. It also serves as a space to add any additional settings as required. For instance, a user could set up an alert to be

generated when there is a change in the orientation of a sensing device (as this is indicated by the accelerometer readings) or when seismic activity is captured by one of the geophones.

6.4. The “Nodes” section

This section shows all the devices currently deployed in a table list format with a focus on the current status of the devices. This view omits recorded environmental data, and only shows technical device information, such as residual energy levels, quality of wireless connection, and so on.

7. Conclusions & Future Work

This work introduced a technological framework for data-driven IoT systems focusing on the application area of landslide monitoring. The three tiers of the framework – namely the data acquisition, data curation and data presentation tiers - were presented in detail and the corresponding considerations and design choices were elaborated. The framework was employed to develop two IoT systems in the context of two pilots for landslide monitoring. The two sites of deployment have landslides of different scales and characteristics and have, therefore, enabled to fine tune and validate the framework in real-life operating conditions. In particular, the use of state of the art IoT technologies has enabled the development of highly scalable and cost efficient systems that provide dense data both in space and time. The two pilots have also been key in demonstrating the added value of data-driven IoT systems in inter-disciplinary research as they nicely complement already existing methods.

This work focused on the methodological approach (as captured by the presented framework) of developing data-driven IoT systems and the technologies employed to develop them. The plan for future work is to allow Pilot 2, at Barton on Sea, to run for two years – such that two full seasons are captured – so as to a) validate and if needed, inform the design of the framework; and b) to collect data such that a dataset of IoT geological data is compiled. This dataset (which will be made available to the community) will be used to train corresponding Machine Learning models,

whose accuracy on forecasting and identifying landslide incidents will be evaluated against currently employed methods. Also, bespoke Visual Analytics will be developed addressing the needs of different stakeholders (scientific community, local authorities, general public). A final strand of our research will be on developing further the framework (in particular the design of the devices) such that their potential environmental impact is reduced or even eliminated.

References

- [1] D. Archer and S. Rahmstorf, *The climate crisis: An introductory guide to climate change*. 2011.
- [2] D. Cruden and D. Varnes, "Landslides: Investigation and mitigation. Chapter 3 - Landslide types and processes.," *Transp. Res. Board Spec. Rep.*, no. 247, 1996.
- [3] D. Varnes, "Slope movement types and processes.," *Spec. Rep. 176 Landslides Anal. Control*, no. Transportation Research Board, Washington, D.C, 1978.
- [4] The International Geotechnical Societies' UNESCO Working Party on World Landslide Inventory, "A suggested method for reporting a landslide," *Bull. Int. Assoc. Eng. Geol.*, vol. 47, no. 1, pp. 5–12, Apr. 1993.
- [5] O. Hungr, S. Leroueil, and L. Picarelli, "The Varnes classification of landslide types, an update," *Landslides*, vol. 11, no. 2, pp. 167–194, Apr. 2014.
- [6] S. L. Gariano and F. Guzzetti, "Landslides in a changing climate," *Earth-Science Rev.*, vol. 162, pp. 227–252, Nov. 2016.
- [7] M. G. Winter *et al.*, "Assessment of the Economic Impacts of Landslides and Other Climate-Driven Events," Dec. 2018.
- [8] British Geological Survey, "East Cliff Landslide, Bournemouth | Landslide case studies," *BGS website*, 2016. [Online]. Available: <https://www.bgs.ac.uk/research/engineeringGeology/shallowGeohazardsAndRisks/landslides/BournemouthLandslide2016.html>. [Accessed: 15-Oct-2019].
- [9] I. West, "Barton and Highcliffe - History and Future of Coast Erosion (Old webpage, restored 2013 with a minor correction in 2019)," 2019.
- [10] Great Britain. Department of the Environment., *Landsliding in Great Britain*. H.M.S.O, 1994.
- [11] T. R. H. Davies and J. F. Shroder, *Landslide hazards, risks, and disasters*, Elsevier. 2015.
- [12] D. M. Cruden and D. J. Varnes, *Landslides: Investigation & Mitigation*. 1996.
- [13] G. Bitelli, M. Dubbini, A. Zanutta, L. Scanning, and R. Sensing, "Terrestrial laser scanning and digital photogrammetry techniques to monitor landslide bodies.," *Interiors*, vol. 35, no. B5, pp. 246–251, 2003.
- [14] N. D. Rose and O. Hungr, "Forecasting potential rock slope failure in open pit mines using the inverse-velocity method - Case examples," in *Proceedings of the 1st Canada-US Rock Mechanics Symposium - Rock Mechanics Meeting Society's Challenges and Demands*, 2007, pp. 308–320.
- [15] X. B. Tu, A. K. L. Kwong, F. C. Dai, L. G. Tham, and H. Min, "Field monitoring of rainfall infiltration in a loess slope and analysis of failure mechanism of rainfall-induced landslides," *Eng. Geol.*, 2009.
- [16] P. R. N. Hobbs, L. D. Jones, M. P. Kirkham, C. V. L. Pennington, D. J. R. Morgan, and C. Dashwood, "Coastal landslide monitoring at Aldbrough, East Riding of Yorkshire, UK," *Q. J. Eng. Geol. Hydrogeol.*, pp. qjegh2018-210, May 2019.
- [17] P. R. N. Hobbs *et al.*, "Establishment of a coastal landslide observatory at Aldbrough, East Riding of Yorkshire, UK," *Q. J. Eng. Geol. Hydrogeol.*, pp. qjegh2018-209, May 2019.
- [18] J. A. Gili, J. Corominas, and J. Rius, "Using Global Positioning System techniques in landslide monitoring," *Eng. Geol.*, vol. 55, no. 3, pp. 167–192, 2000.
- [19] L. Noferini *et al.*, "Using GB-SAR technique to monitor slow moving landslide," *Eng. Geol.*, vol.

Acknowledgements

The authors would like to thank the BCP Council and the New Forest District Council – local authorities of Bournemouth and Barton on Sea correspondingly – for supporting this work. Authors would also like to thank Dr. David Gunn (staff at BGS) for his valuable comments and feedback.

- 95, no. 3–4, pp. 88–98, 2007.
- [20] R. A. Kromer *et al.*, “Automated terrestrial laser scanning with near-real-time change detection – monitoring of the Séchilienne landslide,” *Earth Surf. Dyn.*, vol. 5, no. 2, pp. 293–310, May 2017.
- [21] P. Letortu *et al.*, “Examining high-resolution survey methods for monitoring cliff erosion at an operational scale,” *GIScience Remote Sens.*, vol. 55, no. 4, pp. 457–476, Jul. 2018.
- [22] R. G. L. R. L. B. J. W. K. W. H. S. and L. M. H. Mark E. Reid, “Real-Time Monitoring of Landslides,” *U.S. Geol. Surv. Fact Sheet*, pp. 3008–3012, 2012.
- [23] “Proactive Infrastructure Monitoring and Evaluation (PRIME) System | Geophysical tomography | Waste Management ALERT Project | Our research | British Geological Survey (BGS).” [Online]. Available: <https://www.bgs.ac.uk/research/tomography/PRIME.html>. [Accessed: 04-Feb-2020].
- [24] S. Yinbiao *et al.*, “Internet of Things: Wireless Sensor Networks,” 2014.
- [25] C. M. Angelopoulos, G. Filios, S. Nikolettseas, D. Patroumpa, T. P. Raptis, and K. Veroutis, “A holistic IPv6 test-bed for smart, green buildings,” in *2013 IEEE International Conference on Communications (ICC)*, 2013, pp. 6050–6054.
- [26] A. Herutomo, M. Abdurohman, N. A. Suwastika, S. Prabowo, and C. W. Wijjutomo, “Forest fire detection system reliability test using wireless sensor network and OpenMTC communication platform,” in *2015 3rd International Conference on Information and Communication Technology, ICoICT 2015*, 2015, pp. 87–91.
- [27] W. Gans, A. Alberini, and A. Longo, “Smart meter devices and the effect of feedback on residential electricity consumption: Evidence from a natural experiment in Northern Ireland,” *Energy Econ.*, vol. 36, pp. 729–743, 2013.
- [28] M. V Ramesh, “Real-time wireless sensor network for landslide detection,” *Proc. - 2009 3rd Int. Conf. Sens. Technol. Appl. SENSORCOMM 2009*, pp. 405–409, 2009.
- [29] M. V. Ramesh, “Design, development, and deployment of a wireless sensor network for detection of landslides,” *Ad Hoc Networks*, vol. 13, pp. 2–18, 2014.
- [30] H. hui Wang, X. guo Tuo, G. yu Zhang, and F. ling Peng, “Panzhihua airport landslide (Oct. 3rd 2009) and an emergency monitoring and warning system based on the internet of things,” *J. Mt. Sci.*, vol. 10, no. 5, pp. 873–884, Oct. 2013.
- [31] R. F. Romdhane *et al.*, “Wireless sensors network for landslides prevention,” in *2017 IEEE International Conference on Computational Intelligence and Virtual Environments for Measurement Systems and Applications, CIVEMSA 2017 - Proceedings*, 2017.
- [32] M. El Moulat, O. Debauche, S. Mahmoudi, L. A. Brahim, P. Manneback, and F. Lebeau, “Monitoring System Using Internet of Things for Potential Landslides,” in *Procedia Computer Science*, 2018, vol. 134, pp. 26–34.
- [33] “LoPy.” [Online]. Available: <https://docs.pycom.io/datasheets/development/lopy/>. [Accessed: 05-Feb-2020].
- [34] “DS18B20 Programmable Resolution 1-Wire Digital Thermometer.” [Online]. Available: <https://cdn-shop.adafruit.com/datasheets/DS18B20.pdf>.
- [35] “Soil Moisture Sensor - DFRobot SEN0193 Gravity Analog Capacitive Soil Moisture Sensor.” [Online]. Available: <https://tinyurl.com/rjgeeqq>.
- [36] Sparkfun, “SparkFun Pro RF - LoRa, 915MHz (SAMD21).” [Online]. Available: <https://www.sparkfun.com/products/14916>.
- [37] Atmel, “Atmel SAMD21 Datasheet.” [Online]. Available: <http://ww1.microchip.com/downloads/en/DeviceDoc/SAMD21-Family-DataSheet-DS40001882D.pdf>.
- [38] Sparkfun, “SparkFun 6 Degrees of Freedom Breakout - LSM303C.” [Online]. Available: <https://www.sparkfun.com/products/13303>.
- [39] “Geophone - SM-24.” [Online]. Available: http://www.iongeo.com/virtuals/ResourceArchives/content/documents/Resource Center/Brochures and Data Sheets/Brochures/BR_SEN_Geophones_091509.pdf.
- [40] “ZED-F9P u-blox F9 high precision GNSS module.” [Online]. Available: https://cdn.sparkfun.com/assets/8/3/2/b/8/ZED-F9P_Data_Sheet.pdf.
- [41] Confluent, “Apache Kafka & Event Streaming Platform for the Enterprise, Confluent,”

- confluent.io*. .
- [42] Django Project, “Django,” *Djangoproject.com*. .
- [43] J. Lin, W. Yu, N. Zhang, X. Yang, H. Zhang, and W. Zhao, “A Survey on Internet of Things: Architecture, Enabling Technologies, Security and Privacy, and Applications,” *IEEE Internet Things J.*, vol. 4, no. 5, pp. 1125–1142, 2017.
- [44] M. R. Palattella *et al.*, “Internet of Things in the 5G Era: Enablers, Architecture, and Business Models,” *IEEE J. Sel. Areas Commun.*, vol. 34, no. 3, pp. 510–527, 2016.