Connected Cows: Utilizing Fog and Cloud Analytics toward Data-Driven Decisions for Smart Dairy Farming

Mohit Taneja, Nikita Jalodia, Paul Malone, John Byabazaire, Alan Davy, and Cristian Olariu

Abstract

The Internet of Things (IoT) is about connecting people, processes, data, and things, and is changing the way we monitor and interact with things. An active incorporation of information and communication technology coupled with sophisticated data analytics approaches has the potential to transform some of the oldest industries in the world, including dairy farming. It presents a great opportunity for verticals such as the dairy industry to increase productivity by getting actionable insights to improve farming practices, thereby increasing efficiency and yield. Dairy farms have all the constraints of a modern business — they have a fixed production capacity, a herd to manage, expensive farm labor, and other varied farm-related processes to take care of. In this technology-driven era farmers look for assistance from smart solutions to increase profitability and to help manage their farms well. We present an end-to-end IoT application system with fog assistance and cloud support that analyzes data generated from wearables on cows' feet to detect anomalies in animal behavior that relate to illness such as lameness. The solution leverages behavioral analytics to generate early alerts toward the animals' well being, thus assisting the farmer in livestock monitoring. This in turn also helps in increasing productivity and milk yield by identifying potential diseases early on. The project specializes in detecting lameness in dairy cattle at an early stage, before visible signs appear to the farmer or an animal expert. Our trial results in a real-world smart dairy farm setup, consisting of a dairy herd of 150 cows in Ireland, demonstrate that the designed system delivers a lameness detection alert up to three days in advance of manual observation.

Introduction

The concept of smart dairy farming is no longer just a futuristic concept, and has begun to materialize as different fields such as machine learning have made inroads toward successful applications in this domain. The data-driven approach is transforming many industry sectors including dairy farming, and presents us with an opportunity to predict, control, and prevent certain undesirable events.

The demand for dairy products is rapidly rising due to an ever increasing population coupled with an increase in income per capita [1]. Milk and dairy product consumption is higher in developed countries than in developing nations, but this gap is reducing with increasing incomes, rise in population, urbanization, and dietary changes [2]. It has been estimated that the consumer base of dairy and dairy products is set to rise from from 1.8 billion people in 2009 to 4.9 billion by 2030 [3]. However, methods to improve yield from the agricultural and dairy sector have not advanced at the same rate as the increase in demand. To cope with the increased demand for food, new and effective methods are required to increase the production capacity of this sector. Data-driven decisions, methods, and measures can help increase the production capacity of these industries.

It can be expected that adopting smart dairy farming principles that unify the Internet of Things (IoT), data analytics, fog computing, and cloud computing will help meet these demands and contribute to sustainable growth in the dairy industry. The objective of the work presented is to enable data-driven decisions for dairy farming, and extract timely insights from the data by designing suitable analytics models for such use case scenarios. This aims to provide a set of controls to the farmer and other stakeholders to increase productivity, thus leading to improved farming practices for the overall benefit of the industry.

The rest of the article has been organized as follows. The next section presents the problem space being addressed, Then we present the real-world IoT smart dairy farm testbed deployment, associated challenges, critical decisions, and experience gained throughout the process. Next, we present the design and development methodology used in building the end-to-end IoT solution followed by a technical description of the solution with associated challenges and developed solutions. We then present the benefits to stakeholders, present the conclusion, and discuss ongoing and future work.

The Problem: Early Detection of Lameness in Dairy Cattle

Dairy farmers work hard from dawn until late in the evening, milking, feeding and maintaining the farm. Thus, it is a challenge to monitor the well being of hundreds of cows in a dairy farm in real time. The methods for looking after animal welfare are based on millennia of human experience and grounded in observational methods to analyze animal behavior by visual observation for some kind of anomaly or potential health issue. This leads to the question: *Could technology help? Why can't there be a better way to do it?*

There are behavioral changes when animals become ill, which can be mapped to specific illnesses. The risk of diseases has a large effect on the economy of a farm — payment for veterinary treatments and loss of milk production from the infected animals, as well as animal welfare. *What if one could detect the onset of common diseases before any symptoms are even visible?*

To reiterate, the health and welfare of dairy cows is paramount to the productivity of the herd in both operational and capital expenditure related to pasture management and milk production. One of the issues that need to be addressed in this domain is lameness management.

Lameness is a condition that affects the locomotion patterns of livestock. An all-encompassing definition of lameness includes any abnormality that causes a cow to change the way that she walks, and can be caused by a range of foot and leg conditions triggered by disease, management, or environmental factors. Controlling lameness is a crucial welfare issue, and is increasingly included in welfare assurance schemes.

Lameness is considered to be the third disease of economic importance in dairy cows after reduced fertility and mastitis [4].

Digital Object Identifier: 10.1109/IOTM.0001.1900045

It is estimated [5] that lameness costs an average of €275 in treatment per instance. Early lameness detection allows farmers to intervene earlier, leading to prevention of antibiotic administration and improvement in the milk yield as well as saving on veterinary treatment for their herd.

The existing solutions for lameness detection in dairy cattle either have high initial setup costs and complex equipment, or, in the ones that are technology based, major interoperability issues towards compatibility with existing farm based management solutions. As a solution to this, we have developed an end-to-end IoT application that leverages advanced machine learning and data analytics techniques to monitor the herd in real-time, and identify lame cattle at an early stage.

Real-World Testbed Deployment toward Smart Dairy Farm: Challenges, Decisions, and Experience

Focused on animal welfare and health monitoring, this deployment involves installing sensors on cows' feet. Data generated from these sensors is subjected to analysis using fog computing, which is further enhanced by its cloud component that acts as the site for data fusion and other related resource demanding data analytics functionalities.

The trial was performed on a dairy farm having a herd size of 150 cows in Waterford, Ireland. The important decisions made during deployment and the design phase of the presented IoT solution are listed in this section.

Decision 1: *Which wearable sensor technology should be used from the numerous options available for livestock monitoring?*

From the options available for the sensors/wearables for livestock monitoring, we decided to use the radio-communication-based long-range pedometer (LRP; 433 MHz; industrial, scientific, and medical [ISM] band) instead of a WiFi-based sensor. The reason behind this was that the former does not depend on the Internet for its operation, and serves the purposes of data acquisition in farms where network connectivity is a constraint.

These wearables have lower operational expense and do not use WiFi-based connectivity to send sensed data to a base station. Therefore, as a part of the real-world deployment, off-the-shelf available LRPs (ENGS Systems©®, Israel) specially designed for livestock monitoring were attached to one of the front legs of cows, as shown in Fig. 1. A detailed analysis of other available options and previous approaches were presented in [6].

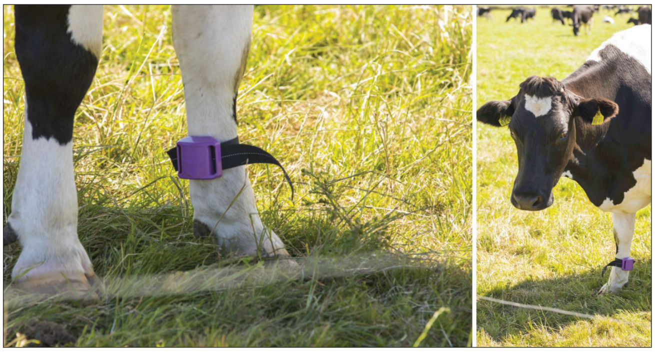

FIGURE 1. Cows with long-range pedometers (LRPs) attached on one of their front legs as part of our smart dairy farm setup.

The workflow and different components of the developed IoT solution are presented in Fig. 2. These pedometers have sampling frequency of 8 ms and forward their sensed data every 6 minutes. The sensed acceleration data is collected at a PC form factor device (fog node), where it is aggregated, pre-processed, and converted into behavioral activities like step count. The system works in both housed and pasture-based dairy systems. The cows are monitored continuously, whether they are in the fields during favorable weather conditions or inside during adverse weather conditions.

In this study, we used three behavioral activities (step count, lying time, swaps) for the analysis with their description as follows:

- Step count: This is the number of steps an animal takes.
- Lying time: This is the number of hours an animal spends lying down, resting.
- Swaps: This is the number of times an animal moves from lying down to standing up.

The choice of these three parameters is based on a literature survey, which suggests these three acts as the best predictors of a lame cow or one transitioning to lameness while analyzing movement or activity patterns of cows.

Decision 2: *Which network device among the available options along the things to cloud continuum should be leveraged as a fog node in such IoT deployments?*

Fog computing is an emerging computation paradigm that aims to extend cloud computing services to the edge of the network, thus enabling computation closer to the source of data. It is being used increasingly in IoT applications, especially in constrained network and Internet connectivity scenarios, which is

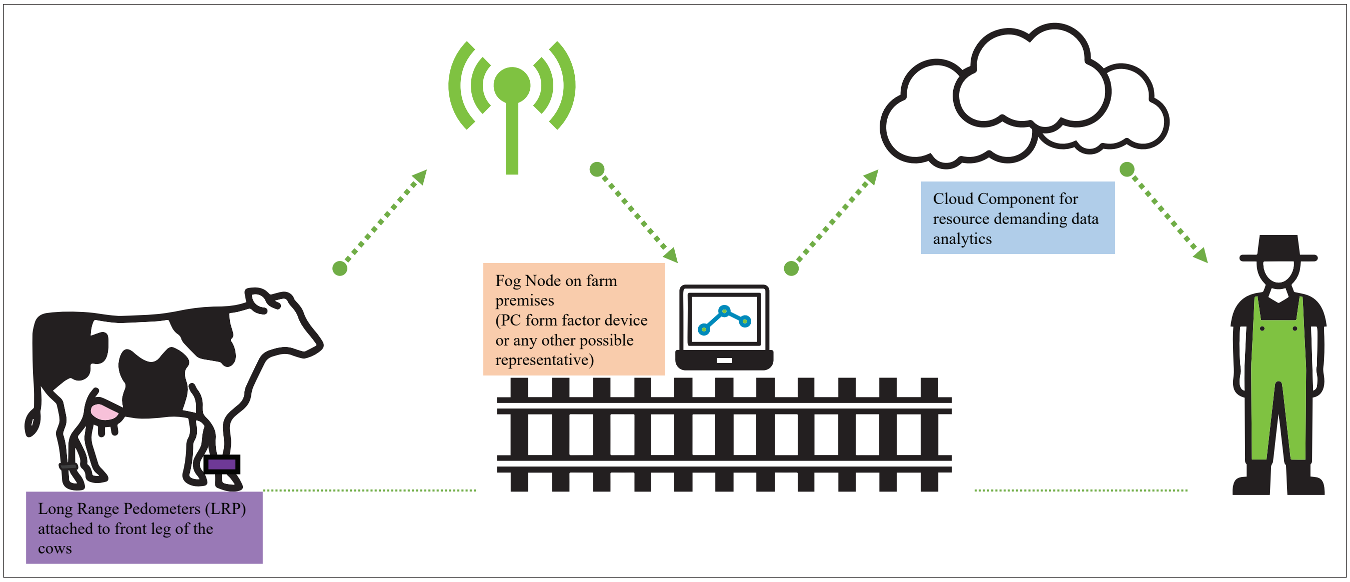

FIGURE 2. System workflow and different components of the developed IoT solution.

also one of the issues in remote farm-based deployments such as ours.

Most IoT-enabled smart farms have some sort of farm management system in place that usually runs on a PC form factor device available within farm premises. Farmers use it to maintain logs and to keep other details electronically at hand. Hence, our plan was to utilize the computing resources already available in such scenarios and leverage them under the fog computing paradigm. Thus, we choose the laptop available with the farmer in our case as the fog node. It should be noted that the developed system is fully able to adapt if the fog node is changed to any other possible representative such as a gateway device. A detailed discussion on this aspect of the system, and also on using resource constrained devices with low computational power as fog nodes, was presented in [6].

This decision also helps to improve fault tolerance, and build up the system resilience to variable farm environments such as weather-based network outages and connectivity issues because of geographically remote locations of farms. In scenarios with low/no Internet connectivity, it becomes ideal to process the data locally as much as possible and send the aggregated or partial outputs over the Internet to the cloud for further enhanced analytical results. The fog-computing-based approach leads to effective utilization of available limited bandwidth and reduces the dependency on the cloud by facilitating part of the data analytics involved in the solution on the network edge. A detailed description of the distribution of services and computational processes running on the edge and in the cloud for the presented solution was described in [6].

Decision 3: *Which streaming protocol do we use for streaming data from the fog component to the cloud component?*

There are a number of options available when it comes to streaming the data, including Message Queue Telemetry Transport (MQTT), Advanced Message Queuing Protocol (AMQP), Extensible Messaging and Presence Protocol (XMPP), and so on. Each of these have their individual pros and cons, and selecting one depends on the use case, objective, and IoT deployment scenario.

Our aim was to use a lightweight protocol that can work in our use case and is also widely supported by both academia and industry in such scenarios. After evaluating and comparing the available options, we selected MQTT as the connectivity protocol in our deployment. It is a lightweight, open source, publish-subscriber model based protocol working on top of the TCP/IP stack, originally invented and developed by IBM [7].

Decision 4: *What should be the development design of the system so that it can be usable, compatible, and able to serve in two user possible scenarios:*
- *When a farmer acts as the end user?*
- *When an agri-tech service provider acts as the end user?*

The end user in our scenario could be a farmer with an existing system or an agri-tech service provider who wants to provide more services to their clients. With that in mind, we decided that the system should be developed as application/software as a service (AaaS/SaaS), which can be used by the service providers to integrate with their existing systems or used directly by the farmer.

This brings us to our next question: *Which software development technique (or architectural style) should be used while developing the system?* The answer and discussion on this is presented in greater detail in the next section.

Designing and Developing Software Systems in Fog Enabled IoT Environments with Cloud Support

Decision 5: *Which software architecture or software development methodology should be used so that the designed system can be multi-vendor interoperable, and also be in line with the finalized design of AaaS/SaaS mentioned above?*

Designing and developing software systems is an intricate process that requires profound understanding of the procedure, consideration of the software architecture and development techniques involved, and knowledge of various interconnected components in the deployed physical or virtual infrastructure.

The microservices architectural style comes as the first realization of a service-oriented architecture and is currently in wide use by industry for software development and deployment as part of best DevOps practices. Given its successful and wide adaptation in the cloud computing domain, a microservices-based architecture seems to be a quite obvious candidate for use in such fog-enabled IoT deployments, but its use is not straightforward. The design and operational practice is sometimes quite different between these two technological paradigms [8]. The major reason for this can be that the microservices approach comes from a different perspective, which is to efficiently build and manage complex software systems, which in turn came to realization as a move toward architectural modularity. The main drivers of modularity are agility, testability, deployability, scalability, and availability

The challenge now is how to apply the microservices approach to build the application in an IoT scenario leveraging the fog computing paradigm. In our analysis, we found that a distributed modular application architecture using microservices was the best approach, given that we could align with the service-based and event-driven needs of our application. Modularity is a must, although not every portion of production has to be a microservice. Microservices need collaboration, and only when there are one or more drivers present should one make use of microservices. In our use case scenario, we had all of the above drivers present. Microservices come with a set of advantages that make it an use an ideal architectural style for software development in end-to-end IoT solutions with constrained environments, giving the ability to overcome the constraints of vendor lock-in, while attributing technological independence between each set of services that make up an application.

Thus, with this understanding we decided on following a hybrid microservices-based approach for application design and development in our end-to-end IoT solution. This decision was also made keeping a future vision in mind of the work, where the microservices act as facilitators to enable dynamic service migration based on the network characteristics to increase quality of service and for better service provisioning.

Technical Challenges and Solutions

Data Analytics and Machine Learning

This section presents details on challenges faced and solutions developed while designing a machine learning model for animal behavior analysis for early lameness detection in dairy cattle.

1) Cow Profiles: *How do we build robust cow profiles that are distinguishable by the learning model as lame and non-lame? Which parameter do we use as a baseline while building and comparing cow profiles?*

For the system to differentiate between normal and anomalous behavior due to lameness, we must first form profiles to characterize normal (non-lame) and lame behavior in the herd. The most frequently used approach for this is to examine the activity level of lame and non-lame animals and study how these differ from the mean of the entire herd. But as it is known that outliers (i.e., a single element in a sample being too high or low) can affect the mean value of a sample, medians or quantiles are sometimes taken as a better measure. To address this issue, we studied the relationship between the herd mean and the herd median. The results of this, as presented in Fig. 3, show that these almost trace out each other for all three activities (lying time, step count, and swaps). This is one of the features of a normal distribution, and therefore it would not matter whether the mean or median is used. Thus, we decided to use herd mean in our analysis.

A study [9] on animal behavior analysis and association patterns of cattle shows that animals grazing within the same pasture can influence the movement, grazing locations, and activities of other animals randomly, with attraction or avoidance; therefore, most of the animals will have their activity levels almost equivalent to the herd mean.

For such reasons, using herd mean as the baseline seems appropriate. Thus, any deviation from the herd mean should serve as a preliminary indicator for a sign of change in behavior, which could potentially be lameness, among other reasons. Such an analysis eliminates the effects of external factors, as these will largely affect the herd as a whole. Further, the measure used to note the deviation in behavior while forming lame and normal profiles of cows in the herd was mean absolute deviation (MAD), while comparing behavior of individual cows with these formed profiles was average deviation.

We build a profile for each animal to characterize normal behavior in a time window using activity-based threshold clustering, details of which are presented later in the article. This helps us to define the lameness activity region (LAR, the period during which the animal is confirmed lame) and normal activity region (NAR, the period during which the animal is confirmed as non-lame), which later acts as ground truth input for the classification model for detecting lameness. An example of this is presented in Fig. 4 for a random cow with ID 2346 in the herd.

However, by comparing the activity of each cow against the herd mean, we found out that not all animals behave the same way. Not all the animals in the herd had their activity tracing the herd mean — some had higher, some lower, and some equal. This observation led us to our next decision in the analysis, which was to identify the clusters in the herd.

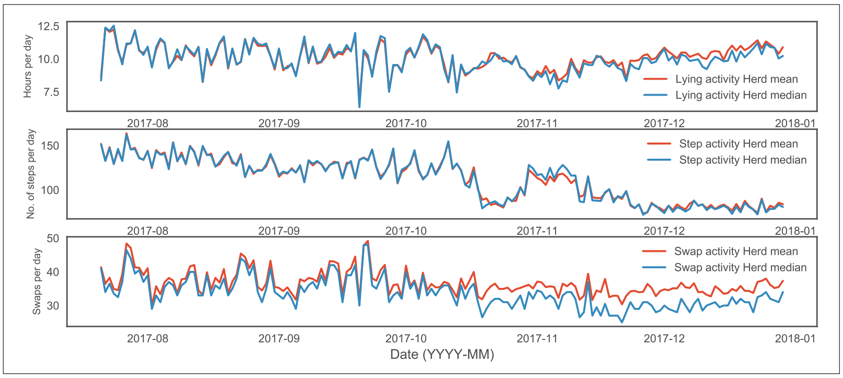

FIGURE 3. Comparing the mean and median of the various animal activities.

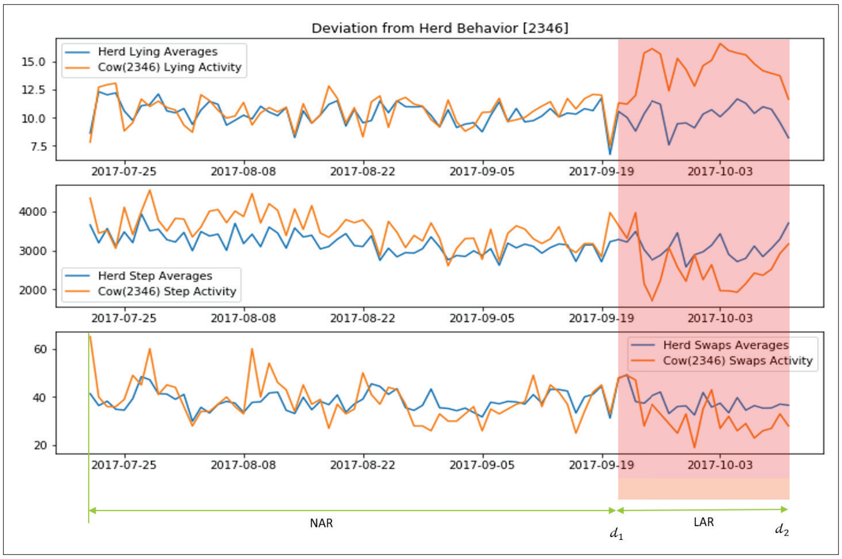

FIGURE 4. Relationship between herd mean and cow activity for cow 2346, showing deviation in its behavior from the herd as it transitions into lameness.

2) Clustering: *Does each animal in the herd need to be treated separately (i.e., treating each cow as a single experimental unit), or can a clustering technique be used to define clusters of animals that share similar features within the herd?*

The same study [9] referred to earlier in forming cow profiles also shows that cattle in the same pasture are not treated as independent experimental units because of the potential confounding effects of the herd's social interactions. It also provides the insight that activity patterns of groups of cows within the herd may have a level of independence that is sufficient for analyzing them as individual units under situations such as large herd size of around 53–240 cows. This means that smaller herds (less than or equal to 40 cows) don't exhibit any patterns of group formations within the herd, while larger herd sizes (53–240) show formations of groups within the herd. It should also be noted that technology-based automated smart solutions for animal welfare are more beneficial for farms with large herds; one can assume that for small ones the farmer can manually keep track of each animal's welfare without much effort.

From our analysis and literature study, it was clear that a one-size-fits-all approach, where it is assumed that all animals behave the same way, and all cows are treated as a single set (i.e. without any grouping) to detect anomaly in behavior, won't be efficient. There are subsets in the herd that share similar features, which once identified can be leveraged to fit the use case as opposed to a one-size-fits-all solution. In our analysis, we found that even animals of the same age behaved differently and had different levels of activity.

Our clustering model is based on the observation that there were some animals in the herd whose activity levels (step count, lying time, and swaps) were always greater than the mean activity value of the herd, and some whose activity levels were always less than the mean herd activity, and then there were others who traced the herd mean. Based on this, we form three clusters as follows:

- **Active:** These are animals in the herd whose activity levels are always higher than the herd mean.
- **Normal:** These are animals in the herd whose activity levels always trace out the herd mean.
- **Dormant:** These are animals whose activity levels are always lower than the herd mean.

It is worth mentioning that prior to finalizing activity-based clustering in our use case, we also used age-based clustering [10] to define clusters and then fed those into the classification model for early detection of lameness. It didn't lead to early detection of lameness, and in line with literature studies we looked for other clustering techniques as well, and found that activity-based clustering performs better [11] in the use case of early detection of lameness.

The above conclusion led to further investigation of the clusters, concerning their nature as static clusters, re-clustering, and optimal approaches to clustering. From our analysis, we found that clusters are dynamic in nature, that is, the animals can migrate from one cluster to another in a time window. There can be a number of reasons behind this; we postulate age and weather at least, and perhaps other factors that affect the activity levels of the animals and the herd as a whole.

Thus, it is the responsibility of the clustering model to re-cluster the animals prior to feeding data into the classification model. The optimal time to re-cluster was found to be about 2 weeks (14 days). This decision was made by continuously observing the movement of animals between different clusters, and finding the time frame of these movements.

3) Classification — Early Lameness Detection: *The next important question was to decide on which classification model should be used given the objective of early detection of lameness in dairy cattle?*

Classification algorithms belong to the set of machine learning algorithms that output a discrete value. Often, these output variables are referred to as labels, classes, or categories. Classification problems with two classes are called binary classification problems, and those with more are referred to as multi-class. In our use case scenario, the problem was written as a binary classification problem, with lame being the positive class and non-lame the negative class. The data split was as 80-20: 80 percent of data was used for model training and the remaining 20 percent was used for testing.

We examined a number of classification algorithms [12] ranging from support vector machine (SVM), Random Forest (RF), K-nearest neighbors (K-NN), and decision trees. We found that the K-NN-based classification algorithm served best for early lameness detection in our use case, as it was best balanced in terms of accuracy and early lameness detection window. It gave an accuracy of 87 percent with a 3-day early prediction window in advance of any visual sign of lameness observed by the farmer.

A short demo video of the overall end-to-end IoT solution thus designed and developed is available at [13].

Benefits to Stakeholders

The detailed impact and benefits to stakeholders are outlined below:

Animals: Animals can't communicate the way humans do. With a little bit of technology, we can understand their natural behavior and trends. We can see the irregularity and change in their behavior and can then take appropriate measures toward their well being. This not only helps improve the production capacity, but it also improves the health and social interactions within the herd.

Farmer: Increased size and scale of the farm poses various challenges for a farmer. In this tech-savvy and data-driven era, it's easier for a farmer to manage the well being of a big herd on a handheld digital device.

Conclusion

We have outlined the key design principles used in the development of our IoT solution aimed at early detection of lameness in dairy cattle. We present the critical decisions made and methodologies used in designing an end-to-end software system in fog-enabled IoT scenarios for our use case.

The key takeaways are:

- A hybrid machine learning model such as the one presented — activity-based clustering combined with a classification model, returns accurate results in detection of anomalies in animal behavior for early detection of lameness as opposed to a one-size-fits-all approach.
- Results clearly suggest that once monitored, the behavioral changes when animals are ill can be mapped to specific illnesses such as lameness in our use case scenario.
- Many of these behavioral changes that occur before visual onset are extremely subtle and difficult to detect in practice without technology.
- A careful coordination of computational resources along the technology path from sensor to cloud continuum is vital to the performance of such a system. Edge, fog, and cloud resources each bring their unique input to the functionality and performance of the overall IoT application system developed.

We believe that the insights from this study can contribute to the behavioral analysis of animals, and can help detect subtle changes in livestock behavior before any clinical symptoms of disease are visible. This will lead to improved insights in animal behavioral analysis and better practices for farmers. The wearable technology for livestock in conjunction with advanced machine learning methods has the potential for development of robust early warning systems to detect disease development early on.

Ongoing and Future Work

To further validate the proposed approach for early lameness detection, we are expanding the work undertaken to date through the execution of a use case in the IoF2020 project (Internet of Food & Farm 2020, https://www.iof2020.eu/) named **M**achine Learning Based **E**arly **L**ameness **D**etection in Beef and Dairy Cattle (MELD). The MELD project is building and expanding on this existing work, integrating it into the IoF2020 dairy farming technology trials with planned deployments in Portugal, Israel, and South Africa, leveraging sensor technologies from two different vendors on a combined total of approximately 1000 cattle. With more data at hand, we then aim to examine other possible clustering techniques and evaluate other classification techniques to further improve the algorithm.

Acknowledgments

This work has emanated from research conducted with the financial support of Science Foundation Ireland (SFI) and is co-funded under the European Regional Development Fund under Grant Number 13/RC/2077. Mohit Taneja is also supported by a Cisco Research Gift Fund. The ongoing work (MELD) is funded through IoF2020, which has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement no. 731884.

References

[1] "EU Agricultural Market Briefs — World Food Consumption Patterns — Trends and Drivers," https://ec.europa.eu/agriculture/sites/agriculture/files/markets-and-prices/market-briefs/pdf/06_en.pdf, June 2015, accessed Dec. 14, 2017.

[2] The Food and Agriculture Organization (FAO) of the United Nations, "Dairy Production and Products: Gateway to Dairy Production and Products," http://www.fao.org/dairy-production-products/en/#.WdTbDmhSyUl, accessed Dec. 14, 2017.

[3] H. Kharas, "The Emerging Middle Class in Developing Countries," OECD Development Centre Working Papers, 2010, http://dx.doi.org/10.1787/5kmmp8lncrns-en.

[4] A. Van Nuffel *et al.*, "Lameness Detection in Dairy Cows — Part 1: How to Distinguish between Non-lame and Lame Cows Based on Differences in Locomotion or Behavior," *Animals*, vol. 5, no. 3, 2015, pp. 838–60; http://www.mdpi.com/2076-2615/5/3/387.

[5] "Economic Cost of Lameness in Irish Dairy Herds," https: //www.xlvets.ie/sites/xlvets.ie/files/press-article-files/XLVets%2520Article%2520Forage%2520Guide%25202012.pdf, accessed May 7, 2019.

[6] M. Taneja *et al.*, "Smartherd Management: A Microservices-Based Fog Computing-Assisted IoT Platform Towards Data-Driven Smart Dairy Farming," *Software: Practice and Experience*, vol. 49, no. 7, 2019, pp. 1055–78, https://onlinelibrary.wiley.com/doi/abs/10.1002/spe.2704.

[7] "Getting to Know MQTT," https://www.ibm.com/developerworks/library/iot-mqtt-why-good-for-iot/index.html, accessed Aug. 3, 2017.

[8] B. Butzin, F. Golatowski, and D. Timmermann, "Microservices Approach for the Internet of Things," 2016 *IEEE 21st Int'l. Conf. Emerging Technologies and Factory Automation*, Sept. 2016, pp. 1–6.

[9] M. B. Stephenson and D. W. Bailey, "Do Movement Patterns of GPS-Tracked Cattle on Extensive Rangelands Suggest Independence Among Individuals?," *Agriculture*, vol. 7, no. 7, 2017; http://www.mdpi.com/2077-0472/7/7/58.

[10] M. Taneja *et al.*, "Fog Assisted Application Support for Animal Behaviour Analysis and Health Monitoring in Dairy Farming," *2018 IEEE 4th World Forum on Internet of Things*, Feb. 2018, pp. 819–24.

[11] J. Byabazaire *et al.*, "Lameness Detection as a Service: Application of Machine Learning to an Internet of Cattle," *2019 16th IEEE Annual Consumer Commun. Networking Conf.*, Jan. 2019, pp. 1–6.

[12] F. Pedregosa *et al.*, "Scikit-learn: Machine Learning in Python," *J. Machine Learning Research*, vol. 12, 2011, pp. 2825–30.

[13] M. Taneja *et al.*, "Smartherd-Connected Cows demo.mp4 — Google Drive," https://drive.google.com/file/d/1QIrKSp8SkAZRRAFQDQHdPXu1T-VYaVyv3/view, shared on Google Drive.

Biographies

Mohit Taneja is currently working as a software research engineer and also a Ph.D. researcher with the Emerging Networks Laboratory in the Telecommunications Software and Systems Group at Waterford Institute of Technology (WIT), Ireland. He is a part of the SFI-CONNECT-IBM project termed SmartHerd, now extended with the follow-on IoF2020 project MELD, which is a project to detect early stage lameness in cattle with the help of technology. His research focuses on fog computing support for Internet of Things applications. He received his Bachelor's degree in computer science and engineering from the LNM Institute of Information Technology, Jaipur, India, in 2015.

Nikita Jalodia is a Ph.D. researcher in the Department of Computing and Mathematics at the Emerg.ing Networks Lab Research Unit in Telecommunications Software and Systems Group, WIT. She is working as a part of the Science Foundation Ireland funded CONNECT Research Centre for Future Networks and Communications, and her research is based in deep learning and neural networks, NFV, fog computing, and IoT. She received her Bachelor's degree in computer science and engineering from the LNM Institute of Information Technology in 2017, with an additional diploma specialization in big data and analytics with IBM.

Paul Malone graduated from WIT with a first class honors degree in 1998 and completed a research M.Sc. in 2001. He has worked on researching technologies and techniques related to IT security in the areas of distributed trust and reputation management and privacy and data protection controls. He is currently coordinating an Agri-Tech use case sub-project of the H2020 IoF2020 EU project applying machine learning technologies in detecting early stage lameness in cattle.

John Byabazaire is currently a Ph.D. student in the School of Computer Science, University College Dublin, Ireland, working on IoT systems for data collection in precision agriculture. Before that, he received an M.Sc. in computer science from WIT. He received his B.Sc. in computer science from Gulu University in 2013. Alongside his current Ph.D. study, he continues to research e-learning for low-bandwidth environments, software defined networking, network function virtualization, remote sensing, IoT, and fog analytics.

Alan Davy received his B.Sc. (with Hons.) degree in applied computing and his Ph.D. degree from WIT in 2002 and 2008, respectively. He is currently head of the Department of Computing and Mathematics in the School of Science and Computing at Waterford Institute of Technology. Previously, he was research unit manager of the Emerging Networks Laboratory with the Telecommunications Software Systems Group of WIT. He is coordinator of a number of national and EU projects such as TERAPOD. His current research interests include virtualized telecom networks, fog and cloud computing, molecular communications, and terahertz communication.

Cristian Olariu received his B.Eng. degree in 2008 from the Faculty of Electronics and Telecommuni.cations, Politehnica University of Timisoara, Romania, and his Ph.D. degree in 2013 WIT in VoIP over wireless and cellular networks. He is currently a research engineer with the Innovation Exchange, IBM Ireland, and previously he was a post-doctoral research fellow with University College Dublin. His interests are in automotive services, IoT, vehicular communications, voice over IP, QoS provisioning for voice over IP, wireless networks, software defined networking, and network function virtualization.