

- ENERGY NEUTRAL MACHINE LEARNING BASED IOT DEVICE FOR PEST DETECTION IN PRECISION AGRICULTURE
- INTERNET OF THINGS AND LoRaWAN-ENABLED FUTURE SMART FARMING
- ADVANCING IOT-BASED SMART IRRIGATION
- PRECISION AQUACULTURE
- CONNECTED COWS: UTILIZING FOG AND CLOUD ANALYTICS TOWARD DATA-DRIVEN DECISIONS FOR SMART DAIRY FARMING
- GROWING PLANTS, RAISING ANIMALS, AND FEEDING COMMUNITIES THROUGH CONNECTED AGRICULTURE: AN IOT CHALLENGE

IoT AND AGRICULTURE



IEEE Access[®]

Multidisciplinary : Rapid Review : Open Access Journal



Become a published author in 4 to 6 weeks.

Get on the fast track to publication with the multidisciplinary open access **journal** worthy of the IEEE.

IEEE journals are trusted, respected, and rank among the most highly cited publications in the industry. IEEE Access is no exception; the journal is included in Scopus, Web of Science, and has an Impact Factor.

Published online only, IEEE Access is ideal for authors who want to quickly announce recent developments, methods, or new products to a global audience.

Publishing in IEEE Access allows you to:

- Submit multidisciplinary articles that do not fit neatly in traditional journals
- Reach millions of global users through the IEEE Xplore[®] digital library with free access to all
- Establish yourself as an industry pioneer by contributing to trending, interdisciplinary topics in one of the Special Sections
- Integrate multimedia and track usage and citation data on each published article
- Connect with your readers through commenting
- Publish without a page limit for **only \$1,750** per article

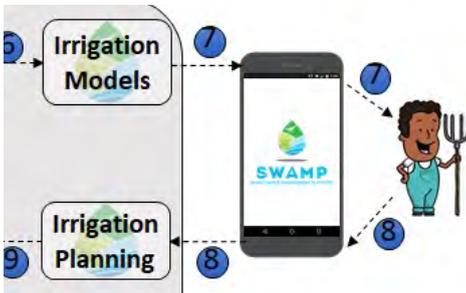


Learn more at:
ieeaccess.ieee.org





Page 10



Page 20



Page 32

2 BRIDGING THE PHYSICAL, THE DIGITAL, AND THE SOCIAL

How IoT and Blockchain Protect Direct-Drinking Water in Schools

6 PRIVACY AND SECURITY

Spectrum Vulnerabilities, Part I: Systematic and Technological Challenges to Identifying and Understanding Vulnerabilities

44 AROUND THE WORLD OF IOT

A Dive into the AgriTech World: Technologies and Adoption Incentives

46 AFTERWORD

IOT AND AGRICULTURE

8 GUEST EDITORIAL

Raffaele Giaffreda

10 ENERGY NEUTRAL MACHINE LEARNING BASED IOT DEVICE FOR PEST DETECTION IN PRECISION AGRICULTURE

Davide Brunelli, Andrea Albanese, Donato d'Acunto, and Matteo Nardello

14 INTERNET OF THINGS AND LORAWAN-ENABLED FUTURE SMART FARMING

Bruno Citoni, Francesco Fioranelli, Muhammad A. Imran, and Qammer H. Abbasi

20 ADVANCING IOT-BASED SMART IRRIGATION

Rodrigo Togneri, Carlos Kamienski, Ramide Dantas, Ronaldo Prati, Attilio Toscano, Juha-Pekka Soininen, and Tullio Salmon Conic

26 PRECISION AQUACULTURE

Fearghal O'Donncha and Jon Grant

32 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

38 GROWING PLANTS, RAISING ANIMALS, AND FEEDING COMMUNITIES THROUGH CONNECTED AGRICULTURE: AN IOT CHALLENGE

Ankita Raturi and Dennis Buckmaster

EDITOR-IN-CHIEF

Keith Gremban, National Telecommunications and Information Administration (NTIA) (USA)

EDITORIAL BOARD

Nuno Carvalho, Universidade de Aveiro (Portugal)
Shawn Chandler, Pacificorp (USA)
Chun-Ting Chou, National Taiwan University (Taiwan)
Robert Douglass, Alta Montes (USA)
Russell Hsing, National Chiao Tung University (Taiwan)
Kaushik Pillalmari, Smartiply, Inc. (USA)
Joern Ploennigs, IBM Research (Ireland)
Arjmand Samuel, Microsoft (USA)
Sachin Seth, Tesla (USA)

Massimo Vecchio, FBK CREATE-NET (Italy)

Tao Zhang, Futurewei (USA)

COLUMN EDITORS

Around the World of IoT

Raffaele Giaffreda, FBK CREATE-NET (Italy)

IoT Standards

Mike Violette, Washington Laboratories (USA)

Policy and Regulatory Issues

Doug Sicker, Carnegie Mellon University (USA)

Privacy and Security

Erin Kenneally, Dept. Homeland Security (USA)

International Computer Science Institute (ICSI)

Scanning the Literature

TBD

Bridging the Physical, the Digital, and the Social
Jun Zhang, University of Denver (USA)

PUBLICATIONS STAFF

Joseph Milizzo, Assistant Publisher
Jennifer Porcello, Production Specialist
Catherine Kemelmacher, Associate Editor
Susan Lange, Digital Production Manager





BRIDGING THE PHYSICAL, THE DIGITAL, AND THE SOCIAL

The scope of the column includes IoT technological achievements that have social impacts and/or incorporate social factors. Each column will provide knowledge and insights in the most recent developments, cutting-edge applications, latest deployments, and conceptual innovations, and of course, their implications on our society. I hope the columns will be meaningful in understanding how our society interacts, adopts, adapts to, and changes with IoT technological advancements.

HOW IoT AND BLOCKCHAIN PROTECT DIRECT-DRINKING WATER IN SCHOOLS

by Zhiguo Shi, Jingxiong Liang, Jun Pan, and Jiming Chen
Zhejiang University

ABSTRACT

The Internet of Things has been regarded as an extension of the Internet and can bring significant changes to our world. A large variety of IoT applications have greatly facilitated our daily lives, such as sharing bicycles, and sharing power banks, for example. These applications optimize resource allocation and thus enhance the efficiency of our society. This article presents a novel IoT application which aims to protect everyday direct-drinking water in schools, via the Internet of Things and blockchain. The system, developed by our IoT team from Zhejiang University and CMCC (China Mobile Communications Group) and deployed in 39 schools in Hangzhou, benefits more than 40,000 students.

INTRODUCTION

Drinking-water safety, especially in schools, has been considered one of the fundamental tasks of the Chinese government. At present, most of the primary and secondary schools in cities in China are equipped with direct-drinking water dispensers with semipermeable membrane filtration. The criterion, "Technical requirements and specifications on the drinking water equipment with membrane treatment for primary and middle schools," issued by the Ministry of Education of China, has been in force since September 1, 2019. The principle of membrane filtration is based on the use of a RO (Reverse Osmosis) membrane to filter out impurities, heavy metal ions, bacteria and other organic pollutants in tap water. The process is implemented by a filter set in the direct-drinking equipment. Usually, these filter sets have a certain life span and need to be replaced regularly according to the volume of the flow filtered. How to effectively evaluate the status of a filter set and get it replaced in time has been an important problem for school managers, as well as for health supervision departments.

To solve this problem, our team devised a high-precision NB-IoT (Narrow Band Internet of Things) water meter with the capability for remote transmission. We also developed automatic processing programs running on our cloud platform to handle these data and send messages to related managers in real time. NB-IoT is one of the most representative IoT communication technologies in the LPWA (low power wide area) category. It perfectly meets the deployment requirements for massive device numbers, low cost, low power consumption, long-term battery lifetime and excellent network coverage. These features greatly facilitate the deployment and maintenance of our system.

On our cloud platform, the interaction between the physical world and the digital world runs as a series of workflows. These workflows are activated by the filter-set replacement detection events generated from the calculated results of the uploaded data, ended by filter-set replacement events submitted by the maintainers or other dispositions made by specific managers with granted permission. By using blockchain, all data are encrypted and stored in blocks of multiple nodes. The system saves a lot of labor costs by shifting the complex management processes in the physical world to automatic processing programs in the digital world. Thanks to the real-time property of the IoT system, the efficiency of all relevant units in supervision has been significantly improved.

PHYSICAL SYSTEM DEPLOYMENT

For direct-drinking water equipment, there are mainly two kinds of maintenance tasks. The first is filter set checking. To monitor the amount of water filtered by a filter set, we deploy our high precision NB-IoT smart water meter in front of the filter set, as shown in Fig. 1. The meter accurately measures the flow purified by the filter set in real time and transmits data to our cloud platform through the NB-IoT network. Basically, every filter set in the equipment has a rated water volume according to its "health permit approval" issued by the local health supervision department. Before reaching the limited volume stated in its "health permit approval," the filter set has to be replaced in advance to avoid degradation to the drinking-water quality.

Another maintenance task is the regular disinfection of the direct-drinking fountains, which is usually done by experienced maintainers before school every day. The supervision and record-keeping of the disinfection process is low efficiency, involving high labor costs and many paper records. In our solution, we propose a way to automatically identify and record disinfection operations by using a mobile app cooperating with low power consumption Bluetooth tags, which are mounted on the equipment to sense and record the information of the disinfection maintainer and operating timestamp. As shown in Figure 2, when the disinfection maintainer stands near the tags, the mobile app will automatically complete the check-in operation and send data to the cloud platform where operating data entry is calculated and stored.

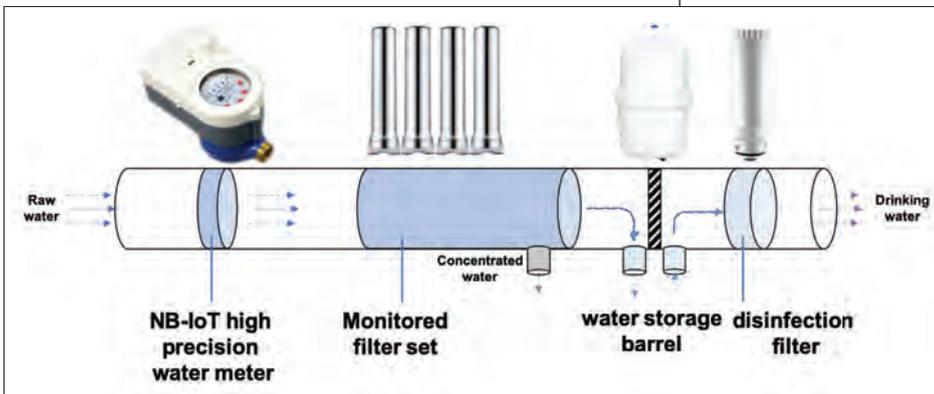


FIGURE 1. Illustration of hardware installation pipeline.

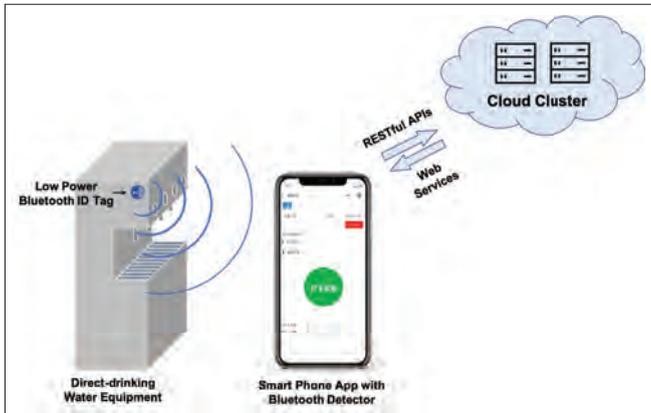


FIGURE 2. Bluetooth check-in scheme.

PLATFORM SERVICES WITH BLOCKCHAIN

When the high-precision NB-IoT water meters collect and upload data to the cloud platform, a series of service programs evaluate the life of the filter set and generate workflow events that send warning messages to the specific maintainer in real time. The replacement of the filter set will be recorded by the maintainer's mobile phone and synchronized to the platform directly. The disinfection operations are processed in the same way. When malpractice is detected, the programs will send warning messages to different superior roles. The platform is accessible to campus administrators, professional maintainers, health supervision departments, and superior management departments. Those departments collaborate efficiently on the platform under different levels of access permissions. The maintenance information of each direct-drinking water equipment is also available to the public through the Internet. Every-

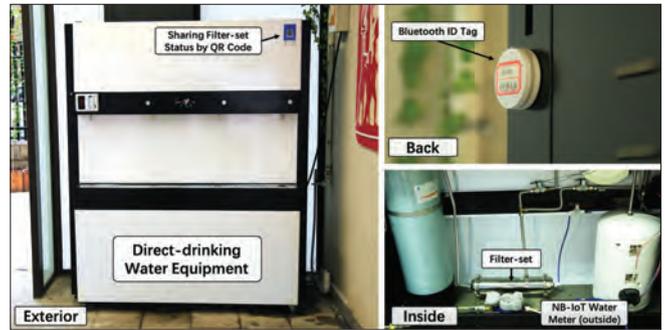


FIGURE 3. Physical system deployment.

one including children's parents can be a supervisor. We also make efforts to improve the transparency and reliability of the supervision data. The IoT system guarantees the data is collected and uploaded in real time. However, there is still a risk for the data to be tampered with, which usually means evasion of responsibility. By using the distributed blockchain network, we have technically ensured that the data cannot be tampered with. Specifically, we built a decentralized blockchain network based on the Hyperledger Fabric, which is the most popular open-source consortium blockchain program at present. On this basis, we encapsulated multiple interfaces in the chaincode for writing and reading these important data and deployed the chaincode on the decentralized network. The properties of the chaincode invocation mechanism and the historical traceability mechanism of ledgers specifically solve the problem of transparent supervision of the direct-drinking water system. Along with real-time transmission, real-time blockchain-based storage guarantees reliability and transparency.



FIGURE 4. Platform panel.

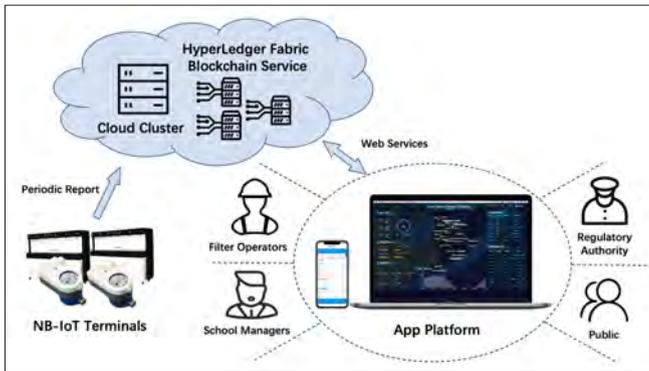


FIGURE 5. System architecture.

SOCIAL EFFECTS

Direct-drinking water in schools usually lacks sufficient supervision. By using IoT and blockchain technologies, we have brought different roles of departments into the regulatory processes, enabling them to collaborate on the platform for efficient supervision. Health-related maintenance information is also made public through the Internet, and the blockchain ensures the credibility of the information. Besides, our platform can also analyze the statistics of water consumption by students through massive equipment data and decide whether to increase the number of equipments in schools. The water consumption of students can also provide some reference for health experts.

SOLUTION PROSPECT

At present, this project has already been put into practice in 39 schools in the Shangcheng District of Hangzhou. There are 326 IoT meters deployed which benefit more than 40,000 students. This project reduces the workload of health supervisors and promotes the traditional on-site supervision to automatic remote monitoring. By using IoT and blockchain technologies, we built this application with reliability and transparency. Now we are working closely with CMCC to bring this project into more scenarios.

BIOGRAPHIES



ZHIGUO SHI [M'10, SM'15] is currently a professor at Zhejiang University and is in charge of this innovation project. He received the B.S. and Ph.D. degrees in electronic engineering from Zhejiang University, Hangzhou, China, in 2001 and 2006, respectively. Since 2006, he has been a faculty member with the Department of Information and Electronic Engineering, Zhejiang University, where he is currently a full professor. From 2011 to 2013, he visited the Broadband Communica-

tions Research Group, University of Waterloo, Waterloo, ON, Canada. His current research interests include signal and data processing, IoT system design, and crowdsensing. He serves as an editor for *IEEE Network*, *IET Communications*, and *KSII Transactions on Internet and Information Systems*.



JINGXIONG LIANG is a major participant in this project. He received the B.Sc. degree from the College of Control Science and Engineering, Zhejiang University, where he is currently pursuing the master's degree. His current research interests include IoT platform and blockchain.



JUN PAN is a major participant in this project. He is currently a Ph.D. candidate in the College of Information Science and Electronic Engineering, Zhejiang University, where he also received the M.Sc. degree. His current research interests include IoT networking, blockchain platforms and smart city systems.



JIMING CHEN [M'08, SM'11, F'19] received the Ph.D. degree in control science and engineering from Zhejiang University in 2005. He is a Changjiang Scholars Chair Professor with the College of Control Science and Engineering, the Vice Dean of the Faculty of Information Technology, and the Deputy Director of the State Key Laboratory of Industrial Control Technology, Zhejiang University. He has published more than 300 papers in premier IEEE transactions/journals and conferences. His research interests include networked control, mobile communications, and connected vehicles, aligned well with the VTS Society.



JUN JASON ZHANG (jun.zhang@du.edu) received his B.E. and M.E. degrees in electrical engineering from Huazhong University of Science and Technology, Wuhan, China, in 2003 and 2005, respectively, and his Ph.D. in electrical engineering from Arizona State University, USA, in 2008. He is currently a professor at Wuhan University, China, and before that he was an associate professor of electrical and computer engineering at the University of Denver, USA. He has authored/coauthored over 100 peer reviewed publications and he is an associate editor of *IEEE Transactions on Computational Social Systems*, an associate editor of *Acta Automatica Sinica*, and Co-Chair of the Blockchain Technical Committee of IEEE SMCS. His research expertise is in the areas of complex systems, artificial intelligence, knowledge automation, blockchain and their applications in intelligent power and energy systems.



JOIN US AT
WF-IoT 2020

IEEE 6TH WORLD FORUM ON THE INTERNET OF THINGS

5-9 April 2020 // New Orleans, LA, USA

The IEEE 6th World Forum on Internet of Things (WF-IoT 2020) is the premier IoT conference of the IEEE Initiatives which will take place in the famous New Orleans on 5-9 April 2020. Every year the world forum is attended by hundreds of most active IoT participants from research community, government and public sector, businesses, multinational corporations and industry.

The technical papers, presentations and events at this conference are focused on contributions to nurture, cultivate, enhance and accelerate the adoption of IoT technologies and applications for the benefit of humanity. WF-IoT 2020 will include a multi-dimensional program of technical research papers, presentations, panels, workshops, tutorials and industry forum on the latest technology developments and innovations in many fields and disciplines that drive the utility and vitality of IoT solutions and applications.

The conference venue will be the Hilton New Orleans Riverside located on the banks of the Mississippi River connected to the Outlet Collection at the Riverwalk, and just a few blocks from the French Quarter.

ABOUT IEEE

IEEE is the world's largest technical professional organization with over 400,000 members dedicated to advancing technology for the benefit of humanity. www.ieee.org

<https://wfiot2020.iot.ieee.org>



PRIVACY AND SECURITY

This column delves into privacy risks of the IoT using risk concepts that are more native to the security domain in order to conceptually bridge our collective understanding, articulation, and management of privacy concerns in the IoT which otherwise might not be sufficiently considered or foreseen by existing legal and technical controls.

SPECTRUM VULNERABILITIES, PART I: SYSTEMATIC AND TECHNOLOGICAL CHALLENGES TO IDENTIFYING AND UNDERSTANDING VULNERABILITIES

by Chris Laughlin

An IoT device in a Faraday cage is about as good as a paperweight. Yet, the inherent openness of wireless systems leaves IoT devices exposed to vulnerabilities. These vulnerabilities are unavoidable precisely because inputs into radio receivers cannot be sealed off and wireless links cannot be completely isolated. Left unaddressed, these spectrum vulnerabilities can be exploited by malicious actors or be the source of non-malicious harmful interference, both of which can have high social and economic costs.

Notwithstanding these vulnerabilities, wireless devices have become increasingly prevalent, and our reliance on radios is only expected to increase.¹ This is unsurprising considering that connected devices have become indispensable to public safety and national security communications, business and critical infrastructure operations, navigation, socializing, and entertainment.

Stakeholders and policymakers have only just begun to understand how the prevalence of connected devices and their vulnerabilities results in a high number of actual and potential security risks. Consequently, they have struggled to identify and understand those risks, let alone develop solutions to address them. If vulnerabilities cannot be identified, they cannot be prioritized, and if they cannot be prioritized, they will not be sufficiently resourced.

In March 2019, I was invited to attend and report on a roundtable discussion convened by Silicon Flatirons that brought together experts from government, academia, and industry to discuss spectrum vulnerabilities.² This column, which will be presented in two parts, captures some of the discussion. Part I will focus on the systematic and technological challenges to identifying and understanding spectrum vulnerabilities. Part II, which will appear in the next issue of *IEEE Internet of Things Magazine*, will focus on recommendations to address those challenges. The ideas and conclusions herein are attributable to the participants, who engaged in the conversation under the Chatham House Rule.³

SYSTEMATIC CHALLENGES

Systematic challenges are those that result from institutional shortcomings, and manifest in inadequate practices, norms, laws, and regulations.

Data Collection and Analysis. Inadequate data collection is one of the most significant barriers to identifying and under-

standing spectrum vulnerabilities. Data collection often fails to occur because harmful interference goes unreported when users assume that equipment is malfunctioning instead of experiencing harmful interference, or because users do not know that such interference should be reported. Even for users that want to monitor interference, it can be difficult to do because devices are not designed to enable such monitoring. At a broader scale, some operators might not be able to detect interference because they outsource expertise about network functions to vendors.

Data Sharing. Even when there are tools to collect and process data, insufficient data sharing among wireless system operators and with the government keeps operators from having information they need to mitigate risks and prevents collaborative efforts to prevent vulnerabilities. Even if there were adequate reporting tools, some operators are reticent to share data because it could contain confidential or proprietary information or information that could be used to assess liability. Additionally, after the National Security Agency's massive surveillance program was revealed in 2013, operators hesitate to share data with the government due to concerns the government would exploit their vulnerabilities.

Research and Testing. Currently, researchers are not getting the data they need to conduct testing. This is in part because data is not being shared with them, but also because academic researchers are increasingly limited by what they can demonstrate in paper and analysis simulations, as a result of inadequate access to real-life networks with which they can experiment. Closed-form solutions are becoming less viable because of increased system complexity, discussed below, not to mention the high cost of proper testing equipment, typically beyond the budget of academic researchers to obtain on their own. Even when research can be conducted, it can be of limited effect, as it often occurs after networks are already commissioned, designed, and deployed.

Government Resources and Authority. Limited government resources and the ways in which those resources are distributed also hinders efforts to identify and address spectrum vulnerabilities. For example, the Federal Communications Commission's Enforcement Bureau, which is limited to civil (and not criminal) enforcement, has fewer than 200 people, split among six divisions, so only a small number cover all the spectrum interference issues across the U.S. Additionally, while there may be dozens of entities in the federal government alone that work in some fashion on spectrum and wireless communications issues, they may not be positioned to collaborate or even know that the other ones exist.

Education. Most policymakers probably cannot explain the difference between 4G and 5G, let alone the diverse and complex causes of spectrum vulnerabilities. As a result, they often do not address vulnerabilities until a harmful event has already occurred. Comprehension of spectrum issues is hardly any better among the general public. Most users do not know when equipment they are using is causing harmful interference or that they might be under the jurisdiction of a federal agency's enforcement authority because the equipment radiates energy.

TECHNOLOGICAL CHALLENGES

Technological challenges are those that result from the design and distribution of wireless systems.

System Design. Equipment commoditization creates misaligned incentives for system developers because it prioritizes qualities like high-speed, low cost, and speed to market over security and resiliency. Even when vulnerabilities in wireless systems are known, they are not given adequate attention. Such is the case for passive intermodulation (PIM), which can be exploited for much more sophisticated attacks,⁴ and 4G LTE specifications, which offer a software stack directly connected to the vulnerable open wireless input. Despite that, those specifications might be used for 5G networks being developed today.

System Complexity. Equipment design vulnerabilities are magnified when vulnerable components are used in the design of other systems. This introduces a level of complexity into a system that obscures the source of the vulnerability because the component's vulnerability can cause insecurity in the entire system with significant consequences. This occurs, for example, when components are purchased because they offer high performance but have low security and yet are being built into devices that have important safety-of-life functions, such as medical devices.

Supply Chain Security and Diversity. The most prominent example of supply chain security risks are concerns that Huawei is building back doors into their equipment, which could later be exploited by the Chinese government.⁵ Equipment with such vulnerabilities could be used, for example, to cause jamming by changing channels within radios to achieve massive pileup pollution within a network. Yet, the equipment is popular for 5G deployments because of its low cost. Over the long term, a lack of equipment diversity has its own risks. The ubiquity of technology from a foreign source could allow an intelligence service to gain leverage, not through a back door, but because it helped build the network, and thereby knows the network better than anybody else does. Equipment diversity risks increase as the number of trusted suppliers decreases due to competition.

Standards Setting. Many participants commented that standards setting is falling short of addressing system design, system complexity, and supply chain issues. System complexity is one reason existing standards are insufficient. Standards setting organizations often develop standards and conduct testing for a particular use, but do not account for the introduction of optional features that are not tested or do not have standards and may have vulnerabilities of their own. Standards are also often developed after innovations are already on the market. Additionally, the U.S. government may be falling short on influencing standards setting bodies, hamstrung by its deference to industry and its cadence of technology neutrality.

Adoption of Standards. Even if the right standards are developed and are effective, system and equipment developers may not implement them. The only teeth the government has to drive standards adoption in the U.S. free market economy is

federal dollars (i.e., procurement and grant money), which is of limited use. For example, consider the case of public safety. There are 330 million people in the U.S., but the public safety population is only about 3 million people, or about one percent,⁶ not to mention the global population and the global market, in which many equipment developers compete. The public safety market does not have enough influence to change industry behavior on its own.

Availability of Harmful Equipment. The risks associated with spectrum vulnerabilities are increasing because there is greater availability of cheap equipment for bad actors. It is only a matter of time before somebody decides to build a jammer to interfere with radio communications for fun. In fact, jammers, despite being illegal in the U.S., can be purchased online. It is difficult for the government to police the entry of malicious, as well as non-compliant, products into the U.S. because of the aforementioned limited government resources.

CONCLUSION

The roundtable identified a number of systematic and technological challenges that prevent identification and understanding of spectrum vulnerabilities. Many, if not all, will need to be overcome to prevent future vulnerabilities and address those that already exist. In Part II of this column, I will discuss many of the solutions proposed by the roundtable participants.

BIOGRAPHY



Chris Laughlin is an attorney in the Communications Practice Group of Kelley Drye & Warren where he represents telecommunications, media, and technology companies in regulatory compliance and policy advocacy before the Federal Communications Commission and on the Hill, litigation in federal district and appellate courts, and transactional matters covering a range of issues, including spectrum access, privacy and data security, and broadband deployment. He earned his J.D. from the University of Colorado Law School and an LL.M. in advocacy from Georgetown University Law Center, where he also served as an attorney and fellow in the Communications and Technology Clinic at the Institute for Public Representation.

FOOTNOTES

- ¹ See Internet of Things Forecast, Ericsson, <https://www.ericsson.com/en/mobility-report/internet-of-things-forecast> (last visited Nov. 18, 2019) (noting there will be 29 billion connected devices by 2022; a small number will be phones and computers using wired connections).
- ² Chris Laughlin, Roundtable Report: Spectrum Vulnerabilities, Silicon Flatirons (June 28, 2019), <https://siliconflatirons.org/publications/roundtable-report-spectrum-vulnerabilities/>.
- ³ Under the Chatham House Rule, participants are free to use and discuss information received, but neither the identity nor the affiliation of any participant may be revealed.
- ⁴ PIM is when the interaction between different physical elements in a network results in interference. See Lou Frenzel, Passive Intermodulation (PIM): What You Need To Know, *Electronic Design* (Mar. 5, 2013), <https://www.electronicdesign.com/wireless/passive-intermodulation-pim-what-you-need-know>.
- ⁵ *Id.*
- ⁶ U.S. and World Population Clock, U.S. Census Bureau, <https://www.census.gov/popclock/> (last visited June 25, 2019).



Raffaele Giaffreda

INTRODUCTION TO THE SPECIAL ISSUE ON IoT AND AGRICULTURE

We are on the verge of a new agricultural revolution. If we look at the extended agricultural sector, intended as the one cultivating plants and farming livestock, with the exception of mechanization and the introduction of chemicals (often referred to as the 3rd agricultural revolution) there have not been too many changes in the past few hundred years.

In fact, due to its intrinsic complexity and variability imposed by type, climate, soil, meteo, etc., looking after a living object (be it a plant, a fish or an animal) has been until now and for the past millennia probably characterized by success practices passed down from father to son. Such a view is supported by statistics like those published by the EU, showing for example that 93.7 percent of all farms are run by only family workers¹: well established and stable business models, best known by the locals and refined over centuries of family trial and error.

Until recent years, challenges posed by climate change as well as globalization are undermining the stability of those models and the status-quo of agricultural businesses as there are many more factors that can influence the successful outcome of cultivations and farming. More meteorological extreme events, global warming, foreign pests and diseases², coupled with a global reduction of arable land³ and an increase of population on earth⁴, all point in one direction: the need for improved quality and quantity of agricultural monitoring data, for more insightful interpretation of cause-effect relationships, and for a more efficient and sustainable use of natural resources such as land, water, etc.

Needless to say, we think IoT technology has a huge role to play in such a landscape, as it can provide an unprecedented source of monitoring data at a very detailed granularity level; with huge amounts of data comes the ability to interpret it for a meaningful and business-viable purpose.

In this Special Issue we touch upon many of these subjects, spanning far and wide between technologies (best picks for wide range connectivity, edge computing, use of machine learning and artificial intelligence) and application domains (from optimized use of irrigation water to fish farming and aquaculture to dairy herds management).

Ensuring wide coverage in rural areas is one of the key enablers to foster innovation in agriculture. As opposed to Smart Cities, rural areas are characterized by customers who

are struggling daily with low margins of running their businesses: covering wide areas for large operators has to be economically viable. Striking such a balance means that low-cost wide coverage can only be guaranteed for so called LPWANs (Low Power Wide Area Network technologies) such as LoRa, SIG-FOX and NB-IoT, to name a few of the most popular ones, all of them characterized by a very low bit-rate supported, per connected customer's device. Having a thin pipe toward the public Internet means that transmitting raw data monitoring is not an option, especially if it consists of images or, even worse, videos. The article entitled "Energy Neutral Machine Learning based IoT Device for Pest Detection in Precision Agriculture" by D. Brunelli *et al.*, focuses on bridging competences on running a low-energy edge computing platform to process data close to the monitoring source and running on it lightweight algorithms trained for the detection of a particular pest, the codling moth, affecting apple orchard cultivations. In this way the LoRaWAN network can be accessed only to communicate a signal if and when the pest is recognized.

The performance of these networks as the numbers of connected devices per gateway grow are also subject of scrutiny. As mentioned above, operators will want to strike the right balance between infrastructure investment and fulfilling the needs of rural communities, knowing that this is a low-margin market where what matters is reaching big numbers thanks to the wide range of their connectivity networks. The article entitled "Internet of Things and LoRaWAN Enabled Future Smart Farming" by B. Citoni *et al.*, after introducing the design details of LoRa and LoRaWAN technology, sheds some light on state-of-the-art achievements and on limitations and bottlenecks of such a technology used in the AgriTech domain.

We already mentioned the need for ensuring that the use of technologies in an agricultural context is economically viable but can also support sustainable practices. The article entitled "Advancing IoT-Based Smart Irrigation" by R. Togneri *et al.* takes a deep cut at one of the most traditional application domains for which we have seen the use of IoT until now: the one supporting smart irrigation. The cost of irrigation water is still not a major concern, but it will not be this way for long given all the effects associated with climate change (droughts and more frequent extreme events are globally reducing the

ability of soil to retain water). But until volumetric water charges for irrigation become widely imposed, adoption will be limited to those contexts where there can be substantial electricity savings from reducing the amounts of irrigation water that need to be pumped. Given the wide variability solving irrigation problems faces, the article proposes a flexible architecture to easily connect IoT and Machine Learning (ML) components to build application solutions in a modular fashion. It shows results from pilot implementations run between Europe and Brazil.

Besides irrigation and precision agriculture, the most popular application domains are where IoT can indeed provide strong support to replace the need for human manpower monitoring, allowing farmers to monitor their farms (fish and livestock) without necessarily being physically present.

In this context, the article entitled "Precision Aquaculture" by F. O' Donncha and J. Grant illustrates how, combining partners' competences and assets from industry, technology and academia, it is possible to provide data-driven insights and decisions that promote ecologically sustainable intensification of aquaculture, taking the example of deployments in a number of fish-farms in eastern Canada.

From Canada to Ireland, moving from fish-farming to connected cows, with the article entitled "Connected Cows: Utilizing Fog and Cloud Analytics Toward Data Driven Decisions for Smart Dairy Farming" by M. Taneja *et al.*, we step into the big issue of being able to monitor the health of cattle just by tracking the animals with some devices (such as pedometers or collars) and recognizing patterns that can be related to a particular condition which, if predicted and controlled early enough, can lead to substantial treatment cost savings. The solution illustrated in this last article is like the others, showing the benefits of being able to monitor environmental conditions through the use of low-cost IoT sensing devices and networks. What one does with collected data in terms of actuation and control and how successful the application can become in terms of busi-

ness, is dependent on the competences of the domain experts to interpret the data and to the ability to relate outcomes to predictions that have a substantial impact on the farmers' bottom line. The more we see of such solutions to farmers' real problems at a sustainable cost, the more we will see the widespread adoption of AgriTech solutions which will accelerate a new data-based revolution in the agricultural sector.

To conclude, I would like to thank the authors and the reviewers for their contributions to this Special Issue and the Editor-in-Chief for the opportunity. It has been an insightful and interesting journey and the hope is that the final outcome will generate the same level of interest for the readers of the IoT community involved at various levels and with different roles in this fascinating application domain.

BIOGRAPHY

RAFFAELE GIAFFREDA (rgiaffreda@fbk.eu) is a chief IoT scientist at FBK CREATE-NET, Italy. He has worked in the telecom R&D environment since the beginning of his career, focusing in the last decade on IoT and related technology transfer activities. In his role, he is now responsible for setting research and innovation directions, acquisition of funding, and the execution of a number of collaborative projects in the IoT domain. He has worked in Italy and the United Kingdom (10 years), acquiring experience in both corporate telco environments (R&D of BT and Telecom Italia) as well as in a small research organization (CREATE-NET before its merger with FBK), where the ability to acquire funding was key to ensuring continuity of operations. He is a recognized expert with a substantial record of IEEE publications and conference presentations, a patent, and various book chapters and tutorials on IoT. He is an experienced speaker and chair of IoT related events, serves as an EU reviewer, and has served on the TPCs of a number of international conferences, and he is the Editor-in-Chief of the *IEEE IoT Newsletter*.

FOOTNOTES

- ¹ https://ec.europa.eu/eurostat/statistics-explained/index.php/Agriculture_statistics_-_family_farming_in_the_EU.
- ² <http://www.fao.org/emergencies/resources/documents/resources-detail/en/c/1180172/>.
- ³ <http://www.fao.org/3/i2280e/i2280e06.pdf>.
- ⁴ https://population.un.org/wpp/Publications/Files/WPP2019_Volume-I_Comprehensive-Tables.pdf.

ENERGY NEUTRAL MACHINE LEARNING BASED IoT DEVICE FOR PEST DETECTION IN PRECISION AGRICULTURE

Davide Brunelli, Andrea Albanese, Donato d'Acunto, and Matteo Nardello

ABSTRACT

Apples are among the topmost fruit crops of the world, and apple orchards are widely expanding in many regions and countries. The most common problem for these crops is the attack of the codling moth, which is a dangerous parasite for apples. IoT sensing devices can nowadays run near sensor machine learning algorithms, thus giving not only the possibility of collecting data over wide coverage but even featuring immediate data analysis and anomaly detection. Near sensor neural network algorithms can automatically detect the codling moth: the system takes a picture of the trap, preprocesses it, crops each insect for classification, and eventually sends a notification to the farmer if any codling moth is detected. The application is developed on a low-energy platform powered by a solar panel of a few hundred square centimeters, realizing an energy autonomous system capable of operating unattended continuously over low power wide area networks. An insightful aspect of this IoT solution is the low power platform for a machine learning algorithm used for IoT fast prototyping. The hardware is based on the Raspberry Pi3 board and the Intel Movidius Neural Compute Stick, responsible for the preprocessing technique and the neural network implementation, respectively. The network model has been analyzed in detail, showing parameter settings and the limitations for the specific hardware constraints. The performance of the proposed system is assessed, and remarks on power consumption are discussed for achieving the zero energy balance of the system.

INTRODUCTION

Recent technological advances have paved the way for remote agricultural sensing and automation. Consequently, sophisticated energy neutral low cost sensors [1] and communication systems [2] can be used as components to monitor and control systems for a sustainable and healthy environment, which is a requirement for smart agriculture applications [3]. However, current wireless sensing platforms and communication systems are designed for bare remote monitoring without making any immediate decision after the damage has already been done [4]. Moreover, the large-scale deployment of sensors would result in a tremendous increase in the number of connections and the amount of data to be transmitted, which could overwhelm current communication systems and also data analysis algorithms [5].

Reconceiving the paradigm of remote sensing operation is imperative to improve the operational performance of precision agriculture. Adding intelligence to the nodes, shifting the detection of anomalies near the sensor to permit decisions and actions as soon as possible, is the key to reduce the communication costs and latencies, and to permit high scalability of IoT solutions in agricultural environments.

Nowadays, machine learning (ML) algorithms are widely used in many fields and are particularly innovative in agriculture to compute tasks such as species recognition [6], water management, crop quality [7], disease detection, and weed detection.

This article focuses on an automatic method for monitoring parasite insects from images taken in pest traps. The codling moth is a particular insect that looks like a butterfly, and it is a dangerous parasite for apple fruit crops. An energy-efficient IoT solution shows how the feasibility of classifying parasites from other general insects autonomously, using low power consumption hardware directly infield. Moreover, the article shows the fast and cost-effective realization of an intelligent sensor and communication system that can be applied in agricultural monitoring and control. It runs ML on the sensor board, and if the insect captured by the camera is classified as a codling moth, a report is sent for an immediate counteraction.



FIGURE 1. Codling moth traps: a) commercial trap; b) prototype of the IoT neural network codling moth smart trap.

IoT SYSTEM

Current methods to monitor pests consist of capturing insects using commercial pheromone-based glue traps, as shown in Fig. 1a, that attract insects even if present at very low densities. Periodic in-field inspections or simple wireless cameras permit the farmer to watch each insect and determine if it is a codling moth [8]. This process is not as smart as an IoT solution could be. In fact, it is slow because it requires the full time presence of an expert, and it is inefficient because even though ML is used, it requires full images sent for remote classification [6].

The proposed system, as shown in Fig. 1b, processes the picture in situ near the sensor (preprocessing algorithm), returns a classification of the insects (ML algorithm) in the trap, and eventually sends a notification to the farmer if it recognizes a codling moth.

As presented in Fig. 2, the system is embedded on a Raspberry Pi 3 that provides the preprocessing stage. Then an Intel Movidius neural compute stick (NCS) with an Intel Myriad X neural accelerator as a vision processing unit (VPU) classifies the images using the model obtained after the training of the deep neural network (DNN).

The system, shown in Fig. 3, has been designed to bring IoT technologies in agriculture where the need to collect the output over vast areas requires long-range communication. Thanks to the onboard intelligence, the output of the smart trap is limited to the few bytes for the report after the classification process,

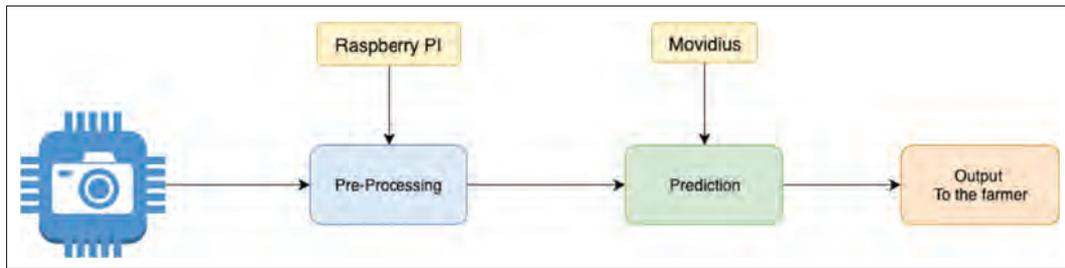


FIGURE 2. Overall system diagram.

and output messages can be managed even with low bit rates. If the farmer needs a visual confirmation from the captured picture, a few images per day can be transmitted as well. Therefore, the trap uses low-power wide area network (LPWAN) technologies and specifically the long-range WAN (LoRaWAN) protocol, which has gained momentum in the market recently. LoRa is a wireless modulation designed for long-range communication at very low energy consumption and bit rate [9]. The LoRaWAN stack defines the communication and security protocols to guarantee interoperability on top of the LoRa network [10, 11].

IMAGE PREPROCESSING AND DEEP LEARNING

The dataset used to start the DNN training contained approximately 1300 pictures and was incremented when more insects were trapped during the earliest experiments. The dataset represents two classes: codling moth and general insects. These figures are used to feed and train the DNN with a TensorFlow model. We used the VGG16 model developed by Oxford University [12]. Then it was converted to a graph model used to perform the classification on the VPU.

The dataset was created with the same camera and trap. The camera captures the bottom side of the insect glue trap; thus, as shown in Fig. 4, pictures may contain a high number of trapped insects to classify. Thus, the images are processed in situ to separate each insect in sub-tiles from the original picture. This step is essential since it filters the raw pictures, as shown in Fig. 4, and produces tiles that contain only one insect. This algorithm is used in two different cases:

- To build a large and comprehensive dataset of pictures for training the DNN model. We started with 100 raw pictures that generated more than 1300 tiles containing only one insect.
- At each application startup, a picture of the trap is taken first, and then, thanks to the preprocessing algorithm, each new trapped insect is cropped for the classification step.

The task efficiently exploits features such as color (dark subjects on a white background) and the shape of the insects with a Blob Extraction algorithm. The process for image crop consists of:

- Conversion of the frame from RGB to gray scale
- Smoothing (or blurring) of the frame with a Gaussian filter
- Edge extraction through a Canny operator
- Some dilation and erosion of the picture

After these operators, the blobs are detected through the OpenCV blob detector. Then each blob is collected in a vector, and the corresponding regions of interest are cropped. All the new pictures are saved for neural network classification.

TRAINING, VALIDATION, AND TEST

For the training stage, we used the rapid development of neural networks for image classification provided by the TensorFlow library [13]. In an ML approach, an initial training step is required. The training consists of an offline process that optimizes the neural network using a large dataset of labeled images. In this way, the system learns the classes assigned to the images. The basic unit of a DNN is the neuron (or node) that multiplies by weight values the input signals. The training phase adjusts the weight values, while some parameters, such as the number of epochs and the image size, can improve the accuracy of a DNN. Epochs represent the number of times all of the training



FIGURE 3. Hardware implementation

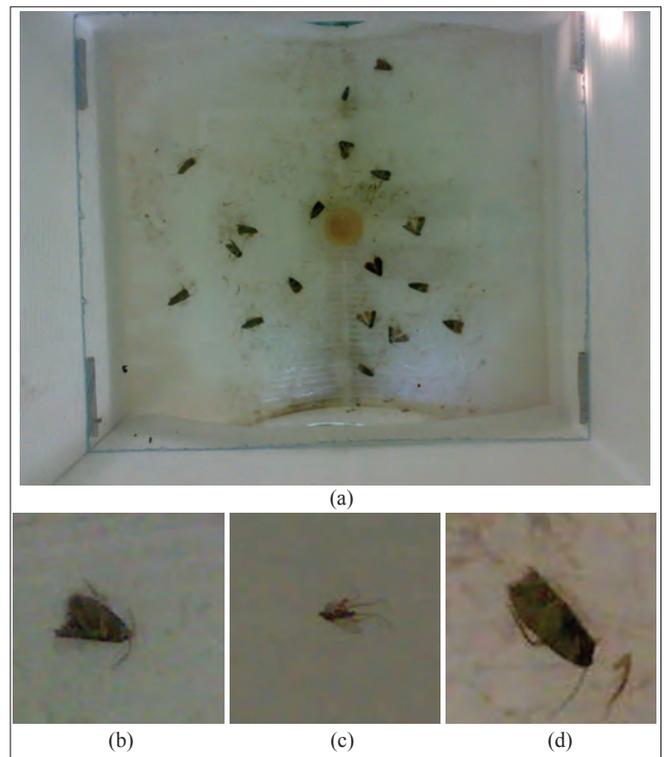


FIGURE 4. Examples of cropped images after pre-processing: a) raw picture; b) tile with a codling moth; c) tile with a general insect; d) tile with a codling moth.

vectors are used once to update the weights. Each epoch finishes with a validation step that evaluates the ongoing training process. A good trade-off between the number of epochs and image size is necessary for a correct training stage and to meet the hardware constraints. The training stage of this application has been assessed with three different configurations:

- 75 epochs, image size 224×224
- 10 epochs, image size 112×112
- 10 epochs, image size 52×52

The results of the training tests are presented in Fig. 5. Notice that training and validation accuracy using 75 epochs

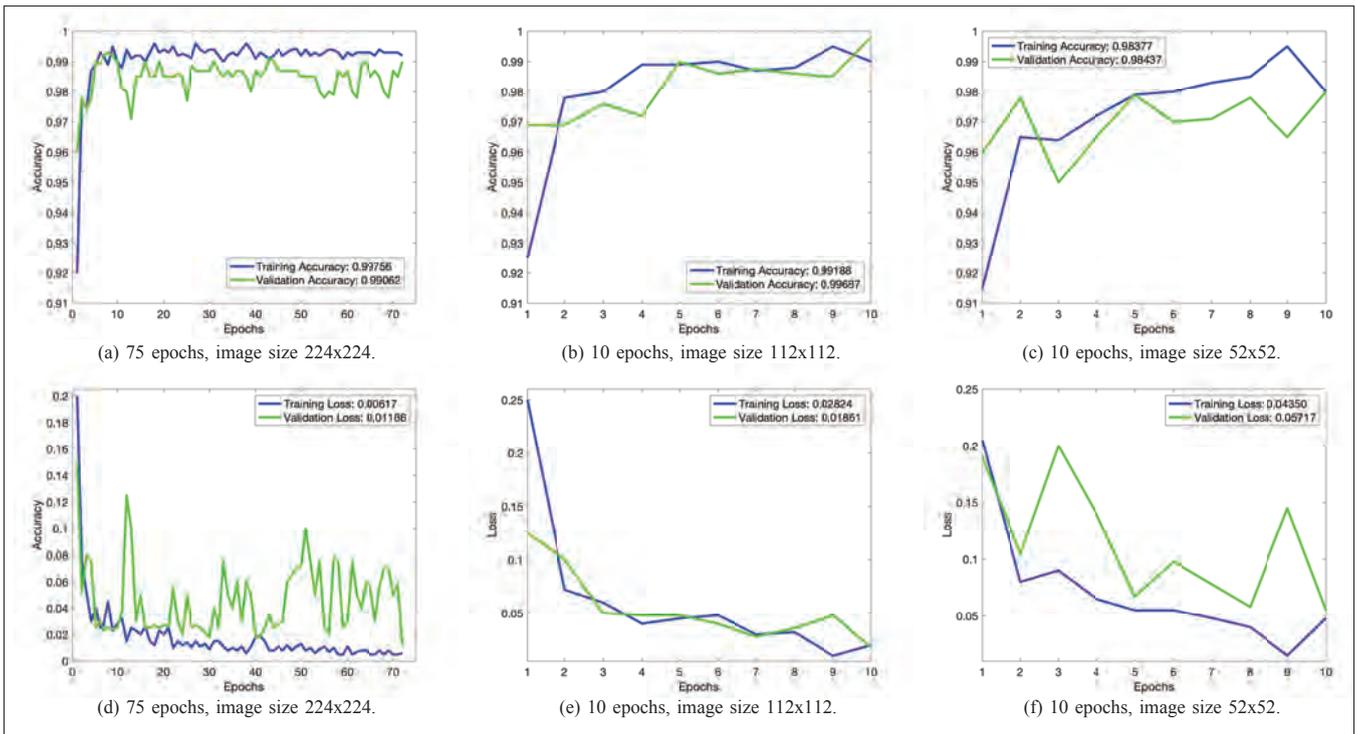


FIGURE 5. Training and validation accuracy and loss function.

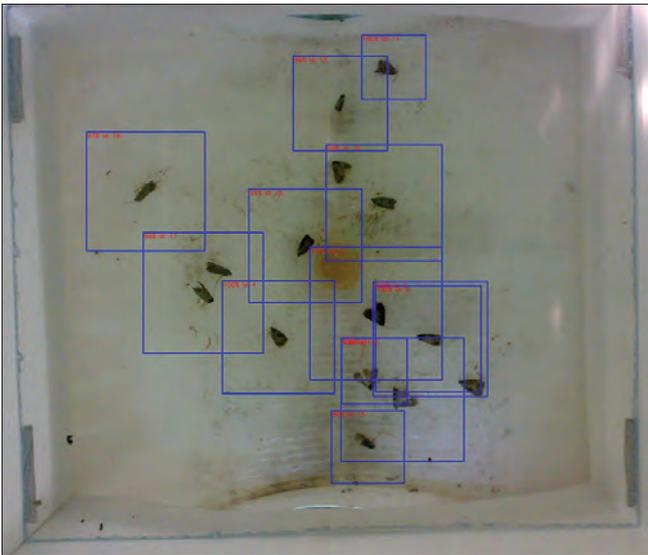


FIGURE 6. Example of moth detection from the system.

is going to be saturated, which suggests that the number of epochs can be decreased to achieve similar performance. As shown in the graphs, 10 epochs are enough for the target accuracy. Moreover, to avoid possible overflow in the Movidius NCS and to save memory on the Raspberry PI 3, the image size can be decreased to meet the hardware constraints because we can use simpler models. We used and tested images of 112×112 and 52×52 pixel size, as shown in Fig. 5. Small images clearly show worse performance with respect to bigger tiles. Nevertheless, about 98 percent accuracy has been achieved, that satisfies the requirements expected by farmers for an IoT service of parasites monitoring.

Figure 6 shows an example of the output from the classification. The DNN provides a confidence measure that indicates how close the detected object is to a general insect or the target Codling Moth.

The tests of the DNN model were carried out during 12 weeks in an apple orchard with the insect glue trap shown in Fig. 1. Tests have involved 262 new insects where:

- 80.6 percent were classified correctly.
- 4.8 percent were false positives.
- 6.4 percent were false negatives.
- 8.2 percent were uncertain.

Thus, the precision is 94.38 percent, the recall is 92.6 percent, and only 8.2 percent need a user assessment watching the raw image.

POWER ASSESSMENT

In apple orchards, codling moth checking is usually executed twice every day. We evaluated the power consumption of the overall system's classification, as shown in Fig. 3, which is divided into five general tasks with different execution time and current consumption:

- Task 0: Boot of the Raspberry (Time 43.68 s, Average Current 345 mA)
- Task 1: Image capture (Time 3.45 s, Average Current 394 mA)
- Task 2: Preprocessing (Time 4.07 s, Average Current 501 mA)
- Task 3: Classification (Time 10.19 s, Average Current 525 mA)
- Task 4: Report/Alarm generation (Time 0.34 s, Average Current 525 mA)

When the system finishes Task 4, it shuts down, and zeroes its power consumption, while a nanowatt real-time clock (RTC) is activated to trigger and boot the application when planned.

As expected, it is possible to observe that T3 is the most power hungry task because it combines the usage of the Raspberry and the Intel Movidius. Figure 7 shows the power consumption of the overall system from T0 to T4, and the total energy necessary is 124.1 J; thus, a 9000 mAh battery is sufficient to sustain the system for more than one year. Moreover, when combining the system with a 0.5 W solar panel of a few hundred square centimeters, as presented in [14], the energy intake will be enough to permit the smart camera to operate unattended indefinitely.

This particular aspect represents a breakthrough for agricultural activities because this means that a farmer could use a smart IoT insect trap, forget about its maintenance, and wait for only automatic alerts if a codling moth is captured.

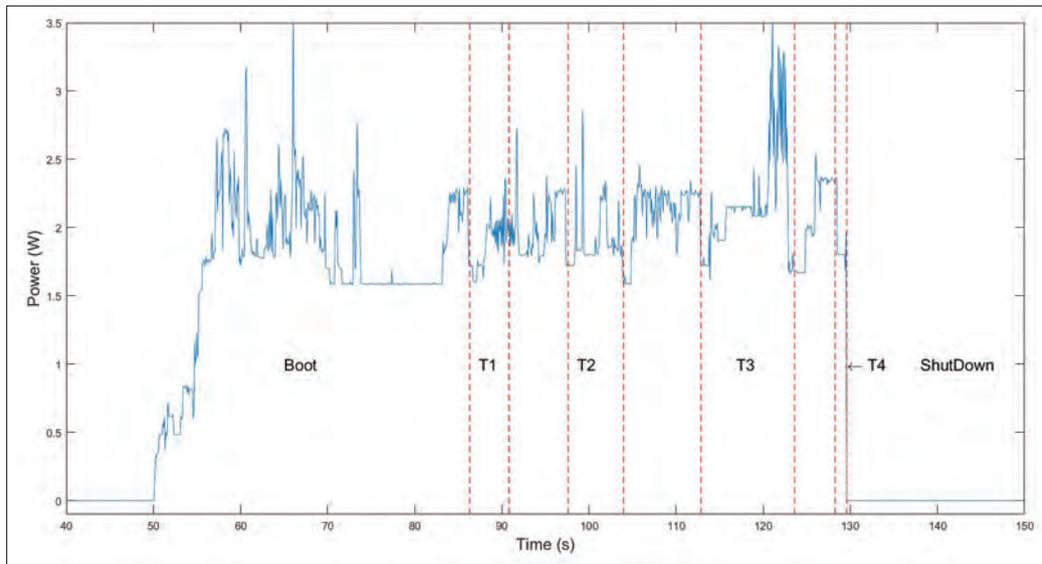


FIGURE 7. System power consumption.

CONCLUSIONS

Even though the proposed system does not use ultra-low-power microprocessors or microcontrollers, its average power consumption is minimal because of its low duty cycle. Due to the low cost of the hardware, this type of system can scale to several installations in the farmer's apple orchard, and save time and money for human intervention in trap checking every day. This type of application is straightforward and innovative, and gives an additional value to agriculture. In this way, it is possible to use treatments for codling moth only when the system detects threats for crops, optimizing the use of chemicals and mitigating their impact on the environment.

ACKNOWLEDGMENT

This research was supported by the IoT Rapid-Proto Labs project, funded by the Erasmus+ Knowledge Alliances program of the European Union (588386-EPP-1-2017-FI-EPPKA2-KA).

REFERENCES

- [1] D. Balsamo *et al.*, "Graceful Performance Modulation for Power-Neutral Transient Computing Systems," *IEEE Trans. Computer-Aided Design of Integrated Circuits and Systems*, vol. 35, no. 5, May 2016, pp. 738–49.
- [2] T. Polonelli, D. Brunelli, and L. Benini, "Slotted ALOHA Overlay on LoRaWAN—A Distributed Synchronization Approach," *2018 IEEE 16th Int'l. Conf. Embedded and Ubiquitous Computing*, Oct 2018, pp. 129–32.
- [3] M. Rossi and D. Brunelli, "Ultra Low Power Wireless Gas Sensor Network for Environmental Monitoring Applications," *2012 IEEE Wksp. Environmental Energy and Structural Monitoring Systems*, Sept. 2012, pp. 75–81.
- [4] D. Sartori and D. Brunelli, "A Smart Sensor for Precision Agriculture Powered by Microbial Fuel Cells," *2016 IEEE Sensors Applications Symp.*, Apr. 2016, pp. 1–6.
- [5] D. Brunelli and C. Caione, "Sparse Recovery Optimization in Wireless Sensor Networks with a Sub-Nyquist Sampling Rate," *Sensors*, vol. 15, no. 7, 2015, pp. 16,654–73.
- [6] W. Ding and G. Taylor, "Automatic Moth Detection from Trap Images for Pest Management," *Comp. Electron. Agric.*, vol. 123, no. C, Apr. 2016, pp. 17–28.
- [7] M. Guermandi *et al.*, "Towards a Wearable Interface for Food Quality Grading through ERP Analysis," *2019 IEEE Int'l. Symp. Circuits and Systems*, May 2019, pp. 1–5.
- [8] A. Guarnieri, G. Molari, and V. Rondelli, "Sistema Elettronico per il Monitoraggio di Insetti Dannosi in Frutteto," *Gestione e controllo dei sistemi agrari e forestali - Memorie AIIA 2011*, 2011.
- [9] L. Tessaro *et al.*, "Lightweight Synchronization Algorithm with Self-Calibration for Industrial LoRa Sensor Networks," *2018 Wksp. Metrology for Industry 4.0 and IoT*, Apr. 2018, pp. 259–63.
- [10] F. Adelantado *et al.*, "Understanding the Limits of LoRaWAN," *IEEE Commun. Mag.*, vol. 55, no. 9, Sept. 2017, pp. 34–40.
- [11] T. Polonelli *et al.*, "Slotted ALOHA on LoRaWAN—Design, Analysis, and Deployment," *Sensors*, vol. 19, no. 4, 2019; <https://www.mdpi.com/1424-8220/19/4/838>

- [12] S. Liu and W. Deng, "Very Deep Convolutional Neural Network Based Image Classification Using Small Training Sample Size," *2015 3rd IAPR Asian Conf. Pattern Recognition*, Nov 2015, pp. 730–34.
- [13] NeuralNetworks; <https://github.com/frank1789/NeuralNetworks>, online; accessed 25 May 2019.
- [14] D. Brunelli *et al.*, "Photovoltaic Scavenging Systems: Modeling and Optimization," *Microelectronics J.*, vol. 40, no. 9, 2009, pp. 1337–44.

BIOGRAPHIES



Davide Brunelli [M'10, SM'16] (davide.brunelli@unitn.it) received his M.S. (cum laude) and Ph.D. degrees in electrical engineering from the University of Bologna, Italy, in 2002 and 2007, respectively. He was a visiting researcher with ETH Zürich, Switzerland, where he was involved in methodologies for energy-harvesting-aware embedded design. He is currently an associate professor at the University of Trento, Italy. His current research interests include smart grids and the development of new techniques of energy scavenging for IoT and embedded systems, the optimization of low-power and low-cost consumer electronics, and the interaction and design issues in embedded personal and wearable devices. He is a member of several Technical Program Committees of conferences in the field of Internet of Things (IoT) and energy management.



Andrea Albanese (andrea.albanese@studenti.unitn.it) received his B.E. degree in electrical and telecommunication engineering from the University of Trento in 2017. He is currently at the same university as a Master's student in mechatronics engineering with a specialization in electronics and robotics. He was involved for three years in the university Formula Student team with a particular focus on power systems and embedded systems. His major areas of interest are machine learning, sensor fusion, and autonomous driving.



Donato d'Acunto (donato.dacunto@studenti.unitn.it) received his Bachelor's degree in electronics and telecommunication engineering at the University of Trento in 2017. He took part in the Formula Student team of the University of Trento for two years working on power management and on the accumulator of an electric car. He is now working on autonomous driving systems and artificial intelligence for surface vessels for his Master's thesis.



Matteo Nardello (matteo.nardello@unitn.it) is currently a Ph.D. candidate at the International Doctoral School in Materials, Mechatronics, and System Engineering at the University of Trento. He obtained his B.S. and M.S. in the field of electronics and telecommunications engineering in 2012 and 2016, respectively. His research interests include the modeling and design of energy-neutral embedded systems, energy harvesting solutions, and long-range low-power wireless communication, with a special focus on autonomous smart IoT devices.

INTERNET OF THINGS AND LoRaWAN-ENABLED FUTURE SMART FARMING

Bruno Citoni, Francesco Fioranelli, Muhammad A. Imran, and Qammer H. Abbasi

ABSTRACT

It is estimated that to keep pace with the predicted population growth over the next decades, agricultural processes involving food production will have to increase their output up to 70 percent by 2050. "Precision" or "smart" agriculture is one way to make sure that these goals for future food supply, stability, and sustainability can be met. Applications such as smart irrigation systems can utilize water more efficiently, optimizing electricity consumption and costs of labor; sensors on plants and soil can optimize the delivery of nutrients and increase yields. To make all this smart farming technology viable, it is important for it to be low-cost and farmer-friendly. Fundamental to this IoT revolution is thus the adoption of low-cost, long-range communication technologies that can easily deal with a large number of connected sensing devices without consuming excessive power. In this article, a review and analysis of currently available long-range wide area network (LoRaWAN)-enabled IoT application for smart agriculture is presented. LoRaWAN limitations and bottlenecks are discussed with particular focus on their effects on agri-tech applications. A brief description of a testbed in development is also given, alongside a review of the future research challenges that this will help to tackle.

INTRODUCTION

World population is expected to grow by 2.3 billion, rising to over 9 billion people before 2050. Most of this growth is expected to happen in developing countries. As a result, the U.N. Food and Agriculture Organization predicts that food production in these countries will need to almost double. Limited and reducing amount of arable land, global climate change, growing scarcity of water, fossil fuel scarcity, and energy price are all factors that will negatively affect the food production process. In this vicious cycle, overpopulation leads to increased demand for agricultural products while also reducing the amount of agricultural destined land, converted into space for infrastructure and housing. Increasing production alone is not enough to achieve food security, and as such we need to look at solutions outside the traditional agriculture methods, creating smarter and technology-enabled agricultural solutions.

The Internet of Things (IoT) is a technological advancement capable of improving efficiency in the global agricultural landscape, accelerating progress toward the goal of increased production. By IoT we mean an architectural framework for systems where computing devices including sensors and actuators wirelessly exchange data collected from everyday objects to either a final user or other machines, in order to monitor and automate processes. It is estimated by market analysts that by 2020 a total of 28 billion devices will be connected to the Internet. This network of Internet connected objects collects relevant data with sensors to be transferred and processed remotely and gives feedback on the current status and actions needed to improve performance. Looking at agricultural applications, these sensors usually gather information about soil and weather condition, animal welfare, crop behavior, and machine status. Smart irrigation systems can, for instance, utilize water more efficiently, only watering the right amount, only in the patch of field which requires it, and at the best time. In turn, this optimizes resource consumption such as water and energy, as well as cost of labor. Such technological advances will play a fundamental role in achieving the prospected increased production requirements when coupled with artificial intelligence (AI)-based programming designed to predict potential issues and provide an adequate response. However, outstanding research issues remain such as developing sensors, communication protocols, and data processing algorithms that can satisfy all the requirements in the context of future smart farming while at the same time being sustainable and cost-effective.

This article is structured as follows. An introduction to low-power wide area networks (LP-WANs) is presented in the next section. Then a review of long-range (LoRa) and Lo-RaWAN is presented, followed by a discussion of current state-of-the-art applications for smart agriculture. We then discuss LoRaWAN technology bottlenecks. Future directions in research are discussed, followed by concluding statements in the final section.

COMMUNICATION PROTOCOLS AND IOT

Fundamental to the IoT revolution is the adoption of a communication technology that can satisfy requirements on three fundamental metrics: energy efficiency, coverage, and scalability. Traditional short-range protocols such as Wi-Fi and Bluetooth, as well as long-range ones such as cellular and satellite communication, fail to provide the required performance to the IoT deployments in smart agriculture and other similar industrial applications. While these protocols in fact are established, neither short-range technologies nor long-range ones are suitable for deployments over a vast area, with sensor nodes that are meant to be "deployed and forgotten": capable of operating for as long as possible with little to no maintenance [1].

Cellular technologies are flawed by design, as they can handle the high data rates of multimedia traffic, allowing only relatively few devices to connect to each base station while granting them wide bandwidth. This is the opposite of what is required by IoT, where a high number of sensor nodes only need the bandwidth necessary to transmit a few bytes every few minutes. This makes long-range technologies impossible to be scaled up without increasing costs and energy consumption. Satellite coverage, while having possibly the best range of all the technologies, is simply too expensive and energy-inefficient for multiple-sensor applications [2]. Short-range communications protocols such as Wi-Fi and Bluetooth suffer partly from the same design flaw as cellular. Although these technologies are in use in some agri-tech applications today, they were also designed to handle a higher volume of data than is required for standard IoT purposes at the expense of increased power consumption, which makes them infeasible for battery-powered devices to be used in rural areas and difficult to access in agricultural and natural environments. These shortcomings were mitigated with the introduction of lower-energy protocols based on IEEE 802.15.4 and designed for wireless sensor networks such as ZigBee. However, their mesh network architecture presents challenges when increasing the amount of connected devices past a certain number without exponentially increasing the network complexity and its power

Digital Object Identifier: 10.1109/IOTM.0001.1900043

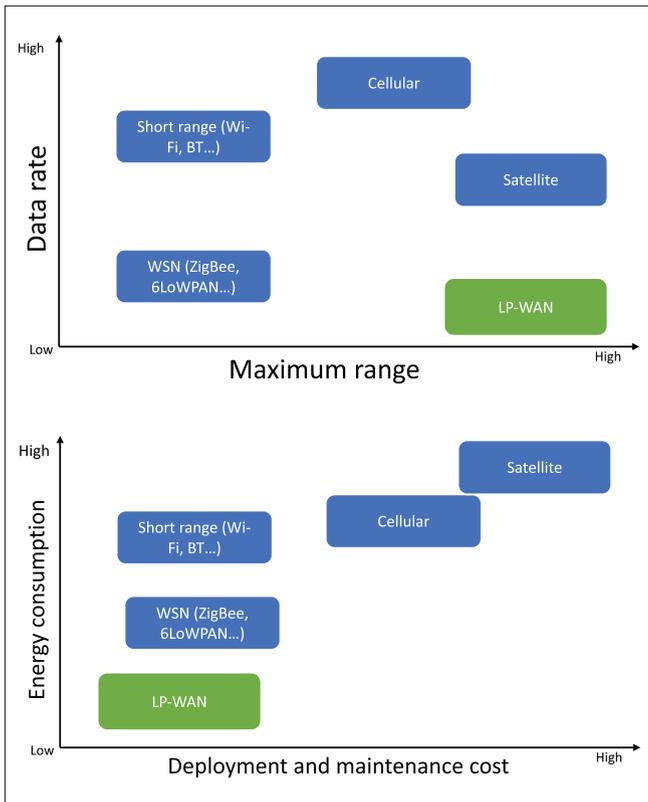


FIGURE 1. Comparison of main IoT enabling communication technologies in terms of range, data rate, energy consumption, and costs.

consumption. For large-scale, sustainable applications over vast areas, both the long-range communication available to cellular and satellite, the lower deployment costs associated with Bluetooth and Wi-Fi and the power consumption of protocols designed with wireless sensor networks in mind such as ZigBee are needed.

LP-WANs are a type of wireless communication network designed to fit all these specifications. While unsuitable for heavy data transmission and multimedia streaming due to very narrow band and low data rate, they can easily support the transmission of small packets of data from sensors and to actuators, minimizing power consumption and design complexity, and thus costs. This is shown in the visual comparison of technologies according to several performance metrics in Fig. 1.

LORA AND LORAWAN

LoRa is a derivative chirp spread spectrum (CSS) modulation technique and proprietary physical (PHY) layer developed by Cycleo to achieve high-range low-power communication. LoRaWAN is a medium access control (MAC) layer protocol built on LoRa, and its specifications are openly available as well as being endorsed by the LoRa Alliance.

LORA PHYSICAL LAYER

LoRa uses the licence-free region-dependent industrial, scientific, and medical (ISM) frequency bands: 863–870 MHz for Europe and 902–928 MHz for the United States. However, it can also be set to operate in the lower ISM bands of 433 MHz and 169 MHz.

While this makes the deployment cheaper due to the use of the unlicensed ISM frequency spectrum, it also restricts the maximum achievable data rate because of regulations on available air time per device on the same frequencies. The enforced duty cycle is 1 percent for the commonly used frequency sub-bands of 863.00–868.00 MHz and 868.00–868.60 MHz. Each

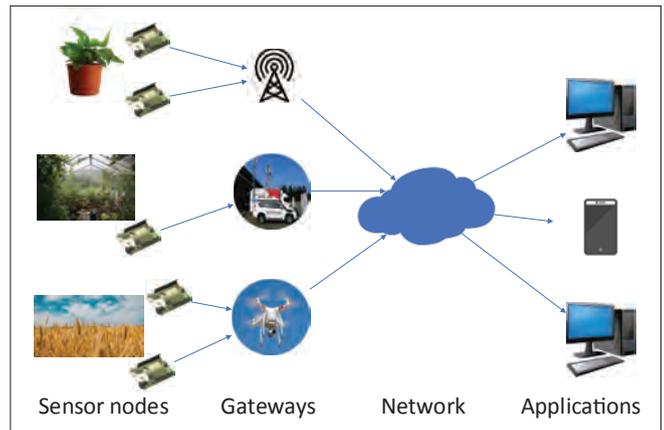


FIGURE 2. Typical star-of-stars LoRaWAN network topology.

sub-band must remain “silent” for a period of time that is proportional to the time on air of the packet and the maximum available duty cycle enforced [3]. For instance, for an air time of 1 s and a duty cycle of 1 percent, the sub-band will have a 99 s mandatory silence time. This restricts the available air time per device to roughly 36 s per day, which makes LoRa unsuitable for high data rate applications. LoRa devices usually can transmit over multiple channels (defined by different center frequencies) and utilize channel-hopping algorithms that aim to find the best possible channel on which to transmit the data in order to try and mitigate this drawback.

Using a lower frequency than Wi-Fi and cellular has the benefit of granting a higher penetration through walls and ultimately a high maximum range. In the range study carried out by J. Petäjärvi *et al.* “On the Coverage of LP-WANs: Range Evaluation and Channel Attenuation Model for LoRa Technology” (2015), this is quoted to be up to 15 km in rural open space and 2–5 km in urban environments, with increased range if there is direct line of sight between devices and gateways.

LoRa has a number of configurable parameters that give flexibility to the designer in regard to the maximum achievable communication range, power consumption, and data rate.

- Spreading factor (SF), which is related to the number of chirps that are used to encode a single bit of information in the modulation of the message. Larger spreading factors increase the signal-to-noise ratio (SNR) and therefore the communication range, at the cost of slower transmission and longer air-time for each packet. Depending on the SF in use, data rate ranges from 0.3 kb/s to 27 kb/s [3].
- Bandwidth (BW), which is the range of frequencies over which the LoRa chirp spreads. Higher bandwidths increase the data rate of packets but reduce communication range. The most common bandwidths available are 125, 250, and 500 kHz.
- Coding rate (CR), which refers to a programmable number of bits that are added to the packet header in order to perform forward error correction techniques. Larger coding rates increase resilience to interference, but also increase packet length, air time, and energy consumption [4].

LORAWAN MAC LAYER

LoRaWAN is one of the available MAC layer protocols built upon LoRa. It has recently gained a lot of attention due to its characteristics, which make it particularly suitable for IoT.

- LoRaWAN networks comprise three main elements:
- Nodes: sensor boards responsible for collecting data or implementing instructions via actuators through communication with gateways
- Gateways: Internet-connected devices that forward the packets coming from the nodes to a network server acting as a logically invisible bridge between nodes and network

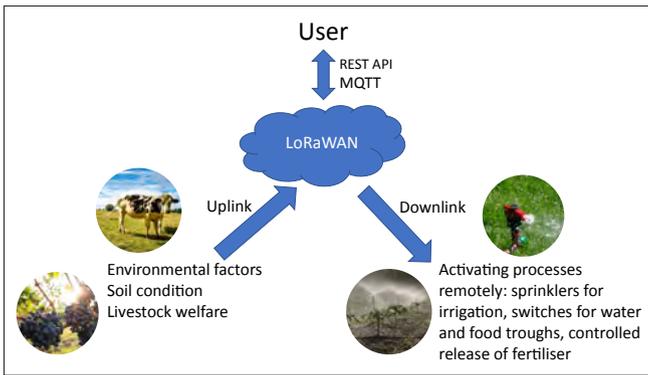


FIGURE 3. Typical downlink and uplink contents in LoRaWAN-enabled agri-tech deployments.

- The network server, which handles the de-duplication of received packets, rejection of corrupted/unwanted ones, as well as scheduling messages to be sent to specific nodes through gateways in range

A typical LoRaWAN network is organized in a star-of-stars topology where nodes do not have a direct connection to any single gateway but instead transmit to all gateways in range. Figure 2 shows an example of topology for a hypothetical smart agriculture application, where different sensing nodes can communicate with different gateways (fixed or mobile based on vehicles or drones). This is different from mesh topologies typically utilized by wireless sensor networks (WSNs), where clusters of devices communicate to sinks they have to associate with directly, and which in turn forward the message on. These multihopping topologies effectively trade off power efficiency for higher transmission range. Thanks to this limitation, the estimated lifetime of a single battery-powered LoRaWAN connected sensor device is expected to be years, which results in cheaper deployment and maintenance as well as an overall simplified network design [1], [2]. The maximum recorded range achieved by an unconfirmed uplink message using LoRaWAN is 702 km.

The gateway relays the data it receives from all the nodes in range to the network server associated with each node. Communication is bidirectional, so devices can send data to the network via uplink and receive instructions via downlink. However, the uplink direction is strongly favored. Direct communication between two nodes is not available with LoRaWAN, requiring the data to pass through a gateway in both uplink and downlink for this to be achieved. These limits preclude the use of LoRaWAN to time-critical, low-latency applications.

Other custom protocols can be built on the LoRa PHY layer. LoRaBlink was developed in [5] to achieve multihop, robust, and low-latency communication, while keeping a low energy profile. Symphony Link™ is another protocol that aims to resolve the problem of scalability outlined by different researchers by implementing a range of different features [6].

LoRaWAN: STATE-OF-THE-ART APPLICATIONS FOR SMART AGRICULTURE

Based on the characteristics outlined so far, LoRaWAN technology has the potential to be applied effectively to many Industrial IoT (IIoT) applications. In these scenarios only small amounts of data need to be analyzed and monitored, with the additional requirement of sending sporadic downlink messages. In an agri-tech context in particular, sensor nodes usually are interested in monitoring environmental factors such as temperature and humidity, as well as health conditions of livestock and chemical conditions of soil and plants. These are often factors that do not require up-to-date real-time monitoring, but need only an update every few minutes. This effectively works around one of the main drawbacks of the LoRaWAN technology, its limitation

in maximum data rate. Downlink messages typically are used to activate simple devices such as solenoid valves and switches like a sprinkler to perform watering of a specific portion of a field or a dispenser to refill food and water troughs for livestock.

The most common IoT applications for agri-tech currently being researched and developed include:

- Automatic irrigation control: optimizing water usage in farming by monitoring soil condition and intelligently activating sprinklers
- Large and small arable farming: including soil monitoring, chemical analysis for pests and disease, machine and agricultural manned and unmanned vehicle (drones) monitoring and control
- Livestock and animal welfare: including movement monitoring to diagnose and prevent diseases such as lameness, eating and drinking habits, and beehive monitoring
- Greenhouse and indoor horticulture: monitoring environmental factors to ensure optimal atmospheric conditions are maintained throughout the year

Generally, academic papers that focus specifically on outdoor, LoRaWAN-enabled agri-tech applications are mostly resolved as proof-of-concept small-scale testbeds for future research, or investigations on the feasibility and performance of the protocol for different use cases. Table 1 includes details of some of the most recent LoRa-specific IoT deployments and some metrics where specified.

On a commercial level, a host of IoT projects have been launched in recent years, mostly in the form of crowd-funded do-it-yourself (DIY) applications ranging from smart gardening gadgets to attempts to automate lawn irrigation. They usually combine LoRaWAN (when used at all) with other technologies such as cellular and Wi-Fi, and are almost exclusively small-scale deployments. In this early stage of the technology, large-scale deployments are still mainly carried out by organizations that can sustain the capital cost of setting up a network as well as providing subscription to servers and data analysis tools run by third parties.

Among the successful examples of such a large-scale LoRaWAN-enabled deployment is the case of livestock monitoring in New Mexico, as reported by Actility, which also exemplifies why LoRaWAN is to be preferred in these scenarios over other network protocols. The amount of cows to monitor (up to 7000) as well as the vast areas these desert ranches occupy (10,000 to 20,000 hectares) makes the process of gathering information about livestock well-being complicated and expensive in terms of time and resources. This is partly due to the amount of animals and area to cover, but also down to stretches of land being only accessible via horse. While historically the cattle was tracked using conventional GPS devices, the lack of consistent cellular coverage over the whole grazing area prevented effective tracking of the cattle's location. A LoRaWAN off-the-shelf solution was able to overcome these problems thanks to its long range and high coverage, while guaranteeing a battery life of 6–7 months to various devices monitoring water level, temperature and GPS position. Where cellular technology failed, LoRaWAN helped increase productivity and security while also reducing the amount of manual work required by business owners.

Moving forward, it is fundamental to try and grow the community of LoRaWAN users as this will lower the costs associated with setting the network infrastructure. As gateways and nodes do not have a 1-to-1 direct connection, a single gateway allows users in its operating range to leverage the established public or private network for their own application. Examples of the benefit this will bring to the community are the case study of Lebanon's Château Kefraya or the Devonian Gardens in Calgary, Canada. Here, thanks to existing nationwide and citywide IoT networks, sensors monitoring, among others, soil temperature and moisture, water temperature, humidity, and luminosity could be set up within a rapid timeframe and with a reduced capital investment [15, 16].

Ref.	Application	Number of nodes	Coverage	SF	BW	Operating frequency	Payload Length	Payload content	Nature of research
[7]	Drip irrigation control	Up to four actuators per node	—	10	125 kHz	433 MHz	Max. 9 B	—	Testbed
[8]	Mushroom house monitoring and control	Three per mushroom house plus actuators	—	—	—	—	—	Temperature, humidity, and CO ₂	Testbed
[9]	Maize crop monitoring	27	648 m ²	—	—	868 MHz	—	Soil moisture and temperature, light intensity, humidity, ambient temperature and CO ₂	Costs and power consumption evaluation
[10]	Tree Farm monitoring	—	Up to 200 m	7 to 11	125, 250 kHz	915 MHz	9 B	Temperature, humidity, solar irradiance, flame sensor	Environmental performance analysis
[11]	Irrigation control	Actuators only	Up to 8 km	12	—	433 MHz	—	—	Proof of concept
[12]	Grape farm monitoring	Three sensor nodes and one actuator node	1 km radius	—	—	—	—	Air temperature, humidity, leaf wetness and soil moisture	Proof of concept
[13]	Water troughs monitoring	Five physical sensor nodes, 100 simulated nodes	0.5 to 2.7 km	—	—	915 MHz	26 B	—	Proof of concept
[14]	Horse stable monitoring	One node	70 m	7	125	868 MHz	2 B	Temperature and humidity	Use case analyses
[14]	Agricultural land monitoring	One sensor buried 10 to 60 cm in soil	40 to 350 m	7 to 10	125	868 MHz	—	Conductivity and soil temperature	Use case analyses

TABLE 1. LoRa-enabled agri-tech applications in the literature.

LORAWAN: LIMITS AND OUTSTANDING RESEARCH LIMITATIONS

Researching the limits of LoRaWAN involves investigating the effects on communication reliability and maximum range upon altering the PHY layer factors of the LoRa protocol: spreading factor, bandwidth and coding rate, which are directly correlated to the time on air of the packet.

Some research indicates that protocol bandwidth configuration has the largest effect on communication range, while other work suggests that the spreading factor choice does instead. We conclude that this debate remains unsolved, while other factors such as temperature, humidity, and antennae position are widely understood to affect communication performance.

Along communication range, packet delivery ratio (PDR) is a fundamental metric to determine how well a sensor node in an IoT network is performing. It is defined as the percentage of packets received over the total amount of packets sent by the end device.

Studies have been published analyzing these two metrics in different environments including urban, mountainous, outdoor and rural, as well as indoor. For these, the data content of the packets is usually not as important as the metadata related to

the packet transmission, which includes information like SNR and received signal strength indicator (RSSI) at the receiving gateway.

In research by S. Wang *et al.*, “Long-Term Performance Studies of a LoRaWAN-Based pm2.5 Application on Campus” (2018), an analysis of air quality on campus ground is carried out, with data sent every 72 s for over 12 months using 22 different nodes with different altitudes and distances from the gateway. Interestingly, devices on rooftops have a lower PDR despite having comparable SNR as devices on lampposts, closer to the ground, suggesting that the SNR is not directly related to the likelihood of a packet being successfully received and decoded. Including additional gateways in the setup, however, results in increased PDR across the whole deployment by about 10 percent.

In “Evaluation of LoRa LPWAN Technology for Remote Health and Well Being Monitoring” by Petäjälä *et al.* (2016), a gateway is placed inside the University of Oulu’s campus and a single moving sensor is set up to transmit a packet every 5 s at +14 dBm. The setup is entirely composed of commercially available products, and results show that using the maximum spreading factor of 12 brings a PDR of around 96 percent. The authors found, however, that the packets were only getting sent

every 13 s instead of the programmed 5. This is because of the limitations imposed by law on the maximum data rate in the unregulated LoRaWAN frequency bands and ultimately will lead to scalability issues.

In fact, as demonstrated by [3], in deployments with 250 to 5000 devices and 3 available channels, not only are the devices constrained to a transmit time that would not exceed the regulations, but also collisions prevent most of the packets from being successfully received and decoded. Because of this, the PDR reduces to values below 20 percent as the number of nodes increase. The problem of collision between packets was also proven by various researchers making use of mathematical models for signal propagation and software simulations.

RESEARCH CHALLENGES

To answer the reduction in PDR that hampers the scalability requirement of any LoRaWAN future application, M. Cattani *et al.* [4] find that it is best to send data using the fastest and more fragile configuration available rather than increasing resilience and air time while trading off speed. Provided that a retransmission function is implemented to handle missed packets and the configuration is such that the deployment exhibits high enough initial PDR (greater than 20 percent), this should yield the maximum effective throughput. On the other hand, in “Performance Analysis of LoRa Radio for Indoor IoT Applications (2017), E. D. Ayele *et al.* carry out indoor performance analysis at the Twente University Campus and reaches the conclusion that the SF should always be increased to minimize the effect of interference and increase PDR across larger distances. Somewhat similar research is brought forward by A. Hoeller *et al.* in “Exploiting Time Diversity of LoRa Networks through Optimum Message Replication” (2018), where each message is sent a number of times, increasing the probability that at least one of those packets is successfully received and decoded by a gateway. This seems to be particularly beneficial for low density networks.

Another avenue of research toward resolving the issue of scalability and packet collision is to investigate the recently deployed adaptive data rate (ADR) mechanism for LoRaWAN v1.1. While its performance has yet to be fully characterized, its goal is to maximize both battery life and network capacity by dynamically altering SF and transmission power of each node. In “EXPLoRa: Extending the performance of LoRa by suitable spreading factor allocations” (2017), F. Cuomo *et al.* present two different algorithms that can assign different spreading factors to the nodes around a single gateway. Somewhat in contrast with the proprietary ADR, which assigns the lowest possible spreading factor that still yields a good communication link between node and gateway, in their work they aim towards a smart and even spreading factors distribution among the nodes. Due to the SF orthogonality, uplink messages sent with different spreading factor can be received by a gateway at the same time on the same channel, hence eliminating a possible collision. The first approach is based on allocating the full range of available spreading factors (7–12) to all sensor nodes. This involves potentially allocating a higher-than-needed SF to some nodes, but by varying the values, the overall probability of collision should decrease. The second approach is an improvement on the first one, taking into account other metrics such as time on air and balancing spreading factors between groups of potential interferes. The algorithms were tested in simulation and showed an overall increase in PDR against the standard ADR algorithm, especially for densely populated networks with fast data rate.

Generally the consensus is that the scalability of LoRaWAN-enabled applications is limited with current state-of-the-art technology [3, 4, 17]. Part of the problem is the fact that downlink availability is itself constrained by the number of nodes a single gateway services. This prevents time-sensitive applications and also limits the possibility for solutions that rely on the feedback of metrics regarding the communication link from the gateway. As there is no node-to-node communication

possible in LoRaWAN, messages between nodes need to be necessarily relayed via a gateway. With gateways subjected to the same duty cycle restrictions as nodes, this could hamper such solutions in a real-life, non-simulated setting.

FUTURE RESEARCH

The following research challenges emerge from the literature reviewed:

- The development of better adaptive data rate mechanisms, based on dynamic spreading factor and other parameters allocation to increase system scalability above the current limits [18]
- The development of retransmission and message duplication mechanisms as an ideal way to deal with collisions and increase PDR as opposed to increasing the spreading factor and time on air [4]
- The reduction of costs and the standardization of hardware and software for LoRaWAN development, which should promote its widespread use

More gateways being online results in an increased downlink capacity for each. This would increase the range of feasible solutions for the issues outlined in this review, all the while reducing collisions [17].

In order to address these issues, several universities and research centers are developing testbeds for development and validation, including our group at the University of Glasgow. The hardware comprises mostly elements purchased through The Things Network (TTN). The Things Network is a community-made website that provides an open source back-end for IoT applications. Three “The Things UNO” nodes are currently set up to monitor air temperature, humidity, and light intensity alongside soil moisture (four variables) for three potted plants, positioned in three different rooms situated on various floors of the University of Glasgow Engineering building, James Watt South. The gateway is also positioned within the building. The vision for this testbed is to expand the number of sensing nodes and gateways to develop and validate different management and data processing algorithms, moving toward adaptive and cognitive implementations that can dynamically self-organize to cope with the network’s changing requirements.

CONCLUSIONS

LoRaWAN has been under the spotlight in recent years due to its suitability to be the standard communication protocol for IoT deployments. It provides long communication range and low energy consumption by drastically reducing the available data rate. In this article, the LoRaWAN protocol was briefly introduced alongside some of the agri-tech applications enabled by it. LoRaWAN’s limitations were also analysed. The biggest issue to future development of large-scale LoRaWAN applications is the effect of packet collision on the deployment scalability. As shown in literature, increasing the number of devices in a deployment with limited gateways drastically reduces the number of packets successfully received and decoded. Duty cycle limitations apply to both sensor nodes and gateways making many of the proposed solutions for packets collision which rely on downlink, such as rescheduling mechanisms or intelligent and dynamic spreading factor allocation, harder to implement or simply not viable.

Many research groups, including the authors’, are working on developing LoRaWAN enabled smart agriculture test beds to improve our understanding of the impact of the presented limitations using experimental test data, and moving towards building predictive models and adaptive network management algorithms for smart farming using the data collected.

ACKNOWLEDGMENT

The authors would like to thank the UK EPSRC supporting Bruno Citoni’s Ph.D. studentship through grant EP/R512266/1, and to acknowledge the support of the wider Communication

REFERENCES

- [1] A. Augustin *et al.*, "A Study of LoRa: Long Range & Low Power Networks for the Internet of Things," *Sensors*, vol. 16, no. 9, 2016.
- [2] R. Sanchez-Iborra and M. D. Cano, "State of the Art in LP-WAN Solutions for Industrial IoT Services," *Sensors*, vol. 16, no. 5, 2016.
- [3] F. Adelantado *et al.*, "Understanding the Limits of LoRaWAN," *IEEE Commun. Mag.*, vol. 55, no. 9, Sept. 2017, pp. 34–40.
- [4] M. Cattani, C. A. Boano, and K. Römer, "An Experimental Evaluation of the Reliability of LoRa Long-Range Low-Power Wireless Communication," *J. Sensor and Actuator Networks*, vol. 6, no. 2, 2017, p. 7.
- [5] M. Bor, J. Vidler, and U. Roedig, "LoRa for the Internet of Things," *Proc. 2016 Int'l. Conf. Embedded Wireless Systems and Networks*, ser. EWSN '16. Junction Publishing, 2016, pp. 361–66; <http://dl.acm.org/citation.cfm?id=2893711.2893802>.
- [6] B. Ray, "Improving LoRaWAN Scalability," 2017; <https://www.link-labs.com/blog/improving-lorawan-scalability>, accessed 31 May, 2019.
- [7] M. Usmonov and F. Gregoretti, "Design and Implementation of a LoRa Based Wireless Control for Drip Irrigation Systems," *2017 2nd Int'l. Conf. Robotics and Automation Engineering*, Dec. 2017, pp. 248–53.
- [8] N. H. N. Ibrahim *et al.*, "IR 4.0 Using IoT and LoRaWAN to Accelerate Lentinula Edodes Growth," in *2018 2nd Int'l. Conf. Smart Sensors and Application*, July 2018, pp. 28–32.
- [9] S. Heble *et al.*, "A Low Power IoT Network for Smart Agriculture," *2018 IEEE 4th World Forum on Internet of Things*, pp. 609–14.
- [10] D. Yim *et al.*, "An Experimental LoRa Performance Evaluation in Tree Farm," *2018 IEEE Sensors Applications Symp.*, 2018, pp. 1–6.
- [11] R. X. W. Zhao *et al.*, "Design and Implementation of Smart Irrigation System Based on LoRa," *2017 IEEE GLOBECOM*, 2017, pp. 1–6.
- [12] D. Davcev *et al.*, "IoT Agriculture System Based on LoRaWAN," *Proc. IEEE Int'l. Wksp. Factory Commun. Systems*, vol. 2018, June 2018, pp. 1–4.
- [13] Lukas, W. A. Tanumihardja, and E. Gunawan, "On the Application of IoT: Monitoring of Troughs Water Level Using WSN," *2015 IEEE Conf. Wireless Sensors*, Aug. 2015, pp. 58–62.
- [14] A. Grunwald, M. Schaarschmidt, and C. Westerkamp, "LoRaWAN in a Rural Context : Use Cases and Opportunities for Agricultural Businesses," 2019, pp. 134–39.
- [15] Actility, "Precision Agriculture"; <https://www.actility.com/precision-agriculture/>, accessed 18 Sept. 2019].
- [16] SEMTECH, "LoRa Technology: Smart Agriculture for Indoor Botanical Garden," <https://www.semtech.com/uploads/technology/LoRa/app-briefs/Semtech-UseCase-Calgary-DevonianGarden-Web.pdf>, accessed 23 Sept. 2019.
- [17] M. C. Bor *et al.*, "Do LoRa Low-Power Wide-Area Networks Scale?" *Proc. 19th ACM Int'l. Conf. Modeling, Analysis and Simulation of Wireless and Mobile Systems*, ser. MSWiM '16. 2016, pp. 59–67; <http://doi.acm.org/10.1145/2988287.2989163>.
- [18] S. Li, U. Raza, and A. Khan, "How Agile is the Adaptive Data Rate Mechanism of LoRaWAN?" *CoRR*, vol. abs/1808.09286, 2018; <http://arxiv.org/abs/1808.09286>.

BIOGRAPHIES



Bruno Cioni received his B.Eng. degree in electronic engineering with music technology systems from the University of York, United Kingdom, in 2016, and his M.Sc. degree in electronic and electrical engineering from the University of Glasgow in 2017. He is currently working toward a Ph.D. in the James Watt School of Engineering at the University of Glasgow, researching IoT, specifically LoRaWAN and its application for smart agriculture and smart cities.



Francesco Fioranelli [SM], CEng, has been an assistant professor at the University of Glasgow since April 2016. He received his Ph.D. degree in electronic engineering from Durham University, United Kingdom, in 2014, and was a research associate at University College London prior to joining the University of Glasgow. He has co-authored more than 60 papers in international conference proceedings and academic journals on various applications of radar systems. His research interests include multistatic and distributed radar sensing, application of machine learning to radar data for classification in healthcare and security, characterization of radar signatures of vehicles and UAVs, wind farms, and sea clutter.



Muhammad Ali Imran [M'03–SM'12] received his M.Sc. (Hons.) degree and Ph.D. degree from Imperial College London, United Kingdom, in 2002 and 2007, respectively. He is the Dean of Glasgow College, UESTC, and a professor of communication systems with the School of Engineering, University of Glasgow. He is an affiliate professor with the University of Oklahoma and a visiting professor with the 5G Innovation Centre. He has university and industry experience, working primarily in the research areas of cellular communication systems. He has been awarded 15 patents and (co)authored over 400 journal and conference publications.



Qammer Hussain Abbasi [SM'16] received B.Sc. and M.Sc. degrees in electronics and telecommunication engineering from the University of Engineering and Technology (UET), Lahore, Pakistan, and his Ph.D. degree in electronic and electrical engineering from Queen Mary University of London (QMUL), United Kingdom, in 2012. In 2012, he was a postdoctoral research assistant with the Antenna and Electromagnetics Group, QMUL. He is currently a lecturer (assistant professor) with the School of Engineering, University of Glasgow. He has contributed to over 250 leading international technical journal and peer reviewed conference papers, and eight books. He received recognition for his research, including appearances on BBC, STV, dawnnews, local and international newspapers, the cover of the *MDPI* journal, most downloaded articles, a U.K. exceptional talent endorsement by the Royal Academy of Engineering, the National Talent Pool Award of Pakistan, the International Young Scientist Award by NSFC China, the URSI Young Scientist Award, a National Interest Waiver from the United States, four best paper awards, and best representative image of an outcome by QNRF. He is an Associate Editor for the *IEEE Journal of Electro Magnetics, RF and Microwaves in Medicine and Biology*, the *IEEE Sensors Journal*, *IEEE Open Access Antennas and Propagation*, and *IEEE Access*, and has acted as a Guest Editor for numerous Special Issues in leading journals.

ADVANCING IoT-BASED SMART IRRIGATION

Rodrigo Togneri, Carlos Kamienski, Ramide Dantas, Ronaldo Prati, Attilio Toscano, Juha-Pekka Soininen, and Tullio Salmon Cinotti

ABSTRACT

Integrated Internet of Things (IoT) platforms are needed for realizing IoT potential in commercial-scale applications. The main challenge is to provide solution flexibility to meet custom application needs. We developed an IoT-based platform for smart irrigation, with a flexible architecture to easily connect IoT and machine learning (ML) components to build application solutions. Our architecture enables multiple and customizable analytical approaches to precision irrigation, making room for the improvement of ML approaches. Impacts on different stakeholders can be anticipated, including IoT professionals, by facilitating system deployment, and farmers, by providing cost reduction and safer crop yields. Examples are given based on pilots in Europe and Brazil.

INTRODUCTION

Nowadays, the Internet of Things (IoT) has already left the state of an idea and has been applied in practical projects. The technical and application challenges are enormous since IoT platforms enable complex real-time control systems that combine the use of communication infrastructure, hardware, software, analytical techniques, and application knowledge combined into multiple layers. One of the key technical challenges is to realize the expected IoT impacts on systems, as IoT allows them to become service mashups, connecting things as services. Consequently, system development will become dynamic plug-and-play interoperable service composition, and system logic will become service orchestration. Overall, IoT allows solution flexibility to fulfill custom application needs.

In the context of agriculture, irrigation is a key task to guarantee adequate crop yield by avoiding under- and over-watering. Moreover, it is an important cost driver, as the energy to transport water and to operate irrigation equipment is costly – in some places, even the water itself is costly. Smart irrigation seeks to apply IoT and analytical methods to leverage precision irrigation, aiming optimal cost effectiveness to the farmer by flowing the water in the proper amount to places where and when it is needed.

In this article we introduce the concept of a flexible IoT-machine learning (ML) platform, wherein IoT and ML components are connected as services in an application context, allowing adaptable solutions to fulfill application needs. This approach benefits IoT professionals, as they can easily develop and deploy complex solutions involving devices, communication, data management, analytics, and application elements.

In particular, our work on this concept has resulted in a platform called SWAMP¹ that implements our flexible IoT-ML architecture toward the smart irrigation problem. This allows highly customizable soil water management solutions, involving flexible connectivity among data, physical models, and ML algorithms oriented to solve application key tasks, such as water need estimation and irrigation planning and operation. We call this concept flexible data-driven soil water management, which in practice allows suitable solutions to a great variety of soil, plant, and regional weather characteristics. This approach benefits the farmer, as a highly customizable smart irrigation solution may reduce water and energy usage and mitigate crop yield risks as it keeps soil water content at healthy levels for plants.

In the remainder of this article, we compare our approach to other practical IoT research projects, provide details on our flexible platform applied to precision irrigation, describe our flexible ML approach to address precision irrigation tasks, highlight the potential impacts of our approach to IoT professionals and farmers, and summarize our main contributions.

RELATED WORK

In recent years, different academic and commercial initiatives have emerged, aiming to incorporate IoT and ML into the agriculture. IoF2020 (www.iof2020.eu)² and Dragon (www.data-dragon.eu) are two projects funded by the European Union for developing IoT platforms for agrifood. IoF2020 is organized into five sectors that adopt different solutions: arable crops, dairy, vegetables, fruits, and meat. Arable crops use sensors to monitor production and intelligent analysis of images to assess crop development. GPS data from cattle neck collars or livestock movements monitor dairy chain, and ML is used for early lameness detection. IoT devices track the production chain of vegetables, fruits, and meat. Dragon aims to integrate IoT data with phenomics, genomics, and metagenomics data associated with ML methods to increase production.

Other recent academic studies include a platform for precision agriculture and experimentation in turmeric cultivation [1]. The platform provides a graphical interface for connecting sensors and actuators and uses analytical methods to analyze the delay of messages. Another study uses thermal images generated by drones and transmitted throughout a cloud-fog system to identify non-uniform irrigation zones [2].

Commercial companies are also putting some of these ideas into the market. Examples include Agrosmart (www.agrosmart.com.br) in Brazil, Agricolus (www.agricolus.com) in Italy, and Cropmetrics (www.cropmetrics.com) in the United States. Among other technologies, they use soil sensors, weather stations, and weather forecasts for irrigation advising. However, their underlying approaches for water need estimation and irrigation optimization and operation are not available.

Despite the richness of recent approaches adopting IoT in the agri-food chain, they do not yet fully explore: (a) the architectural aspects to hold flexible solutions involving IoT and ML components and (b) the potential of the collected data for a more accurate analysis of water needs. In this article, we advocate that data-intensive methods provided by ML algorithms, in combination with IoT technologies and weather-soil-atmosphere simulations, can provide a considerable impact in precision irrigation and its solution deployment.

A FLEXIBLE IoT-ML PLATFORM FOR SMART IRRIGATION

The success of next generation systems for precision irrigation based on IoT technologies coupled with intelligent ML data processing techniques depends on the ability of the solution to adapt to different contexts found in farms. A flexible IoT-ML platform must allow different deployment configurations of hardware, software, and communication technologies, customized to deal with the requirements and constraints of different settings, countries, climate, soils, and crops. Here we advocate that an IoT-ML infrastructure for providing smart irrigation services be defined by two complementary dimensions,

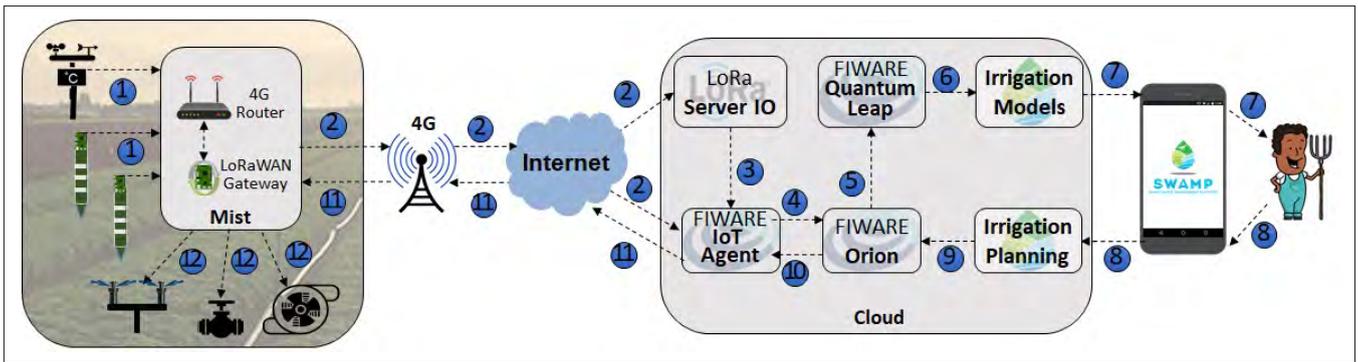


FIGURE 1. IoT-ML platform overview and end-to-end data path for a smart irrigation scenario.

namely core components and deployment locations. Core components are a set of software, hardware, and communication technologies, such as soil moisture sensor probes, long-range WAN (LoRaWAN) [3], Message Queuing Telemetry Transport (MQTT) [4], LoRa Server (www.loraserver.io), FIWARE [5], and SEPA [6], as well as specific services for water need estimation and irrigation planning and operation.

On the other hand, deployment locations define where core components can be placed, and how they communicate with each other. This generates distinctive configuration scenarios for different deployments. Locations follow an IoT computing continuum, composed of things (sensors and actuators), mist (field nodes such as radio gateways), fog (farm on-premise computing infrastructure), cloud (data storage and processing), and terminal (a smartphone, tablet, or laptop where the end user interacts with the application). The five instances of this continuum define the end-to-end information path starting with data collected by sensors up to commands executed by actuators. The five instances might not necessarily be present in all scenarios. Rather, depending on farm characteristics, requirements, and constraints, fog or cloud may not be present. This feature provides additional flexibility to the IoT-ML platform, as the differences are understood, and the platform adapts to the farm and not the opposite.

Figure 1 depicts the IoT infrastructure for providing smart irrigation services, composed of core components and deployment locations. In smart agriculture, each farm has particular objectives and characteristics, so different deployment configurations may be used, representing instances of the same platform. Figure 1 presents a simpler version of the deployment of the IoT-ML platform where locations are thing, mist, cloud, and terminal (i.e., no fog is used). This configuration was chosen for simplicity and a farmer's choice of not hosting any on-premises infrastructure.

In Fig. 1, the numbers in blue circles represent a simplified sequence of the end-to-end data flow through this deployment of the IoT-ML platform. Soil moisture sensors send data via LoRaWAN to the gateway installed in the mist node. Particularly for the SWAMP Project, we have built a custom-made three-depth soil moisture sensor, but also use commercial sensors from Libelium (www.libelium.com) and Meter (www.metergroup.com). A weather station also sends data to the mist node via a serial wired interface (1). From there, the mist node forwards data via 4G through the Internet directly to the cloud (2).

Within the cloud, sensor data are treated by the LoRaWAN server and sent to the IoT protocol translator (3), such as a FIWARE IoT Agent. Weather data goes directly to the IoT protocol translator using the Ultralight 2.0 protocol, as well as weather forecasts obtained from an external service. The Translator converts the three different types of input data — soil, weather conditions, and weather forecast — into the format of the particular IoT underlying platform transmitting them to the context broker (4) (e.g., NGSi JSON format for FIWARE

Orion). Once data arrives at the context broker, it is forwarded to time series storage (5) that makes it available for further processing (e.g., FIWARE QuantumLeap using CrateDB), where the first part of the end-to-end data flow ends. Here, depending on the volume and velocity of data, the time series storage may be replaced by a distributed data pipeline (e.g., Apache Kafka) connected to a big data processing system (e.g., Apache Spark). However, for most smart agriculture scenarios dealing with individual farms with hundreds of sensors, time series storage is a lightweight solution that provides adequate performance.

The water need estimation component obtains soil moisture, weather conditions, and weather forecast data from the time series storage (6) to generate ideal crop water need estimates. Water need estimation is further divided into physical and ML models, further explained in the next section. The estimates are in turn used by irrigation planning to generate an optimized and real plan that is aware of different physical and financial constraints (7). Farmers are shown the irrigation plan via the Farmer App (8) and approve or change the irrigation plan that is sent back to irrigation operation (9), which controls the irrigation system. From there, irrigation commands follow the way back to the mist going through the context broker (10), IoT protocol translator (11), and Internet/4G (12). Finally, irrigation commands reach sprinklers, pumps, and valves (13).

Should the fog be present in a different scenario, the end-to-end communication would be preserved with small changes, as some components would be deployed on-premises in the farm office where the fog node is located, such as the LoRaWAN server. This scenario includes the direct operation of the irrigation system. In alternative scenarios, the irrigation plan either interacts with existing third-party irrigation systems (e.g., Netafim — www.netafim.com — or Focking — www.fockink.ind.br) already installed on the farm or even used by farmers to operate the irrigation systems manually. These options for interacting with an irrigation system are common, and we assume they are generic enough to represent an IoT ecosystem for smart irrigation.

FLEXIBLE DATA-DRIVEN SOIL WATER MANAGEMENT

Over the last decades, data-driven soil water management has been accomplished using physical models³ or triggering soil moisture sensors data⁴ [7]. However, as IoT enables more abundant data, with major spatial and temporal granularity and low latency, it increasingly allows the rise of data-driven approaches. Here, a key challenge is how to make all analytical techniques (e.g., physical models and ML) available to work together in a flexible IoT platform, considering that each crop, type of soil, and region may demand a different solution. In this sense, our work aims to show the roles that these analytical techniques can play in soil water management, and how they can be flexibly assembled together. Our approach is based on two main characteristics:

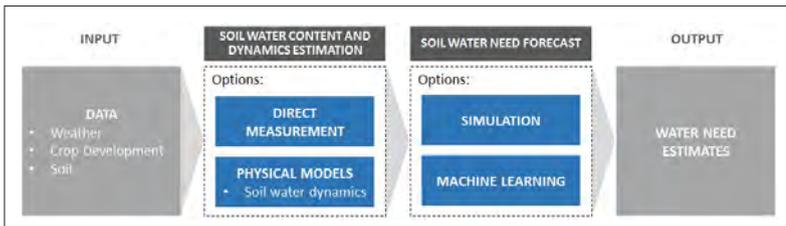


FIGURE 2. Water need estimation process.

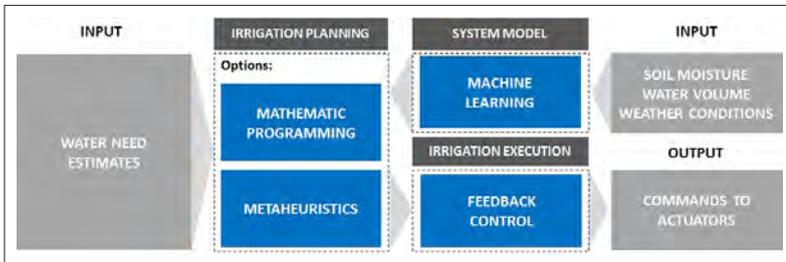


FIGURE 3. General data flow between irrigation services.

- Solution flexibility: Modularized components are integrated as services in the IoT platform, allowing solution flexibility.
- Increased ML relevance: Analytical solutions can be assembled from traditional physical approaches, and modern ML and simulation ones, and even by combining them. Benefits of this approach are not only improved irrigation plans, but also self-improved platforms that learn from experience.

Soil water management is performed in two phases, namely water need estimation and irrigation planning.

WATER NEED ESTIMATION

The precision irrigation problem can be modeled as the soil water balance system at the root zone, where the soil water content is the result of the balance between water content level and a series of mechanisms that make this level increase or decrease (soil water dynamics) [8]. IoT provides the ability to monitor water content levels and dynamics, while soil water management systems seek to maintain water content in an optimal range [7].

Figure 2 depicts the water need estimation process, divided into two key activities:

- Soil water content and dynamics estimation: This consists of estimating soil water content and dynamics through:
 - Direct measurement of soil water content, rainfall, irrigation, and so on
 - Physical models of soil water dynamics applied over the collected data (weather data mainly) and soil and crop characteristics
- Soil water need forecast: This consists of calculating soil water content forecasts and water need forecasts for each moment of a planning horizon, using techniques such as simulation and ML algorithms:
 - Simulation is appropriated when working with physical models, iteratively applied to simulate future data points [9].
 - ML takes advantage of data from multiple time series (soil moisture, soil water balance, soil characteristics, and weather data) of direct sensor readings or from variables derived out of physical models. Weather forecast data, provided by external services, can also be used. Different multivariate forecasting methods can be used to handle these multiple time series [10, 11].

The process depicted in Fig. 2 allows flexibility once it allows different components combinations, as they are implemented as services in the platform. It is also possible to customize each combination, as they have numerous options inside them (i.e., different physical models and ML techniques).

As an example, the SWAMP project provides two customized analytical solutions among all combination possibilities to fit the characteristics and needs of different pilots. One of them, called CRITERIA-1D [9], uses physical models and simulation, wherein soil physical dynamics models (physical models) are the input to soil simulation that generates soil water content and water need forecasts. Another solution uses direct measurements, physical models, ML, and simulation, wherein the main input is direct measurement of soil moisture enriched by an evapotranspiration⁵ model (physical model) [7] as the main soil water balance contributor. ML techniques⁶, such as Panel VAR [10] and RNN-LSTM [11], are used for the processing of soil water content and water need forecasts, and simulation to test alternative irrigation scenarios. Note that the latter uses Panel VAR and RNN-LSTM, respectively, a traditional and a cutting-edge technique for time series, thus highlighting the solution flexibility in exploring different ML techniques as they gain relevance.

IRRIGATION PLANNING AND OPERATION

The water need estimation models provide what can be called the ideal irrigation. There are, however, other aspects that need to be considered when conceiving an actual irrigation plan, that is, a plan that can be put in place in the farm, which include:

- Water availability: Water scarcity is a problem in various parts of the planet. Water quotas or supply schedules might not allow the ideal amount of water to be irrigated in time. If the needed amount of water is not available, the irrigation plan should allocate the existing water so that the best economic return to the farmer is achieved.
- Costs of irrigation: Even if the water comes from private reservoirs, irrigation is not free. Pumping the water to the fields consumes energy, and its cost has an impact on the farmer's bottom line. For example, in certain regions of Brazil, the energy bill can account for up to 30 percent of the production cost. A cost-aware plan should avoid irrigation when tariffs are higher.
- Limitation of the irrigation systems: Irrigation methods differ in how much of the irrigated water actually reaches the plants: furrow irrigation has 60 percent efficiency, while sprinkler irrigation reaches 75 percent [12]. Other aspects of the irrigation infrastructure need to be considered when planning: maximum pumping capacity of the farm, uniformity of irrigation, and soil variability, among others.

Figure 3 presents a modular approach that separates irrigation planning from operation, completing the data flow shown in Fig. 2. There are three main modules:

- Irrigation planning: Computes the timing and water volume of irrigation events that best address the crop needs, while being aware of operational constraints and economic interests. Linear and nonlinear programming techniques can be used, as well as approximate solutions such as those provided by metaheuristics [13].
- Irrigation operation: Communicates with the sensors and actuators installed in the farms, sending commands and monitoring the operation to ensure adherence to the plan. It controls the opening and closing of valves, the pressure at pumps, and so on, using the underlying IoT communication infrastructure to send commands. The use of standard IoT interfaces and protocols enables on-demand addition of sensors and actuators, smoothing the transition toward fully automated irrigation.

- System Model: Computes an updated model of the system behavior as far as irrigation is concerned. IoT devices (e.g., soil sensors, water meters) in combination with data-driven techniques enable estimating the actual irrigation efficiency, and planning accordingly.

DISCUSSION AND LESSONS LEARNED

PHYSICAL MODELS VS. MACHINE LEARNING

The increasing use of IoT in precision irrigation brings spatial and temporal accuracy gain, as sensors can potentially be placed in all manageable locations on a farm. Thus, site-specific particularities can easily be considered, leveraging the exploration of data-driven approaches. In this context, the question of what would be the right combination of techniques to deliver adequate soil water management emerges. Do traditional approaches, such as using physical models or triggering soil moisture sensors data [7], still take place? Or is this the time to avoid physical models and use cutting-edge ML algorithms acting directly to data?

Our vision is that there is no unique ideal analytical approach for all cases, as crops differ significantly in irrigation methods and in crop, soil, and regional weather characteristics. More than that, depending on crop culture or region, not all data features might be available or cost-effective. However, a discussion of the roles traditional and ML approaches can play in effective solutions can provide guidelines to discern the most appropriate alternatives to each application case.

Physical models have been extensively used in irrigation, bringing implicit agronomic knowledge, as they connect raw data features to specific and relevant features. Nevertheless, there are important limitations, as general models involve simplifications that often ignore local particularities, while site-specific models work well only regionally, and few models have adequate performance levels for different regions. Finally, the few models of general application that are flexible enough to address different conditions [7] are often complex and require many data features that are difficult to obtain.

On the other hand, pure ML approaches applied directly to IoT data seem promising, as cutting-edge deep learning is capable of capturing implicit knowledge from raw data in many application areas, as well as delivering highly customizable results [14]. For this reason, we believe that ML approaches will be extensively explored in scientific research in the coming years, allowing the emergence of truly cognitive smart irrigation systems. As such, our architecture approach, based on core components and deployment locations, gives the necessary flexibility not only to build customizable IoT-ML solutions, but also to assemble customizable data-driven solutions. For smart irrigation, we have shown that it is possible to use various combinations of analytical tools, including mixes of physical models, simulation (traditional approaches), and ML techniques.

As ML gains momentum, existing physical models may lose room because ML could implicitly capture from raw data the same information physical models provide. Instead, IoT's continuous growth might enhance the utilization of physical models, as they can calculate their outputs with better spatial and temporal granularity. Also, IoT tends to promote not only physical models but also ML. In summary, a futuristic vision may be that ML is well positioned for IoT-based applications. However, although ML seems to have a promising future for smart irrigation, we are still at the beginning of its exploration, and we need reliable data that now is still generated by physical models or a combination of data and ML techniques.

All in all, considering the advantages and disadvantages of each side, we advocate that current solutions consider using both physical models and ML algorithms — physical models serving as feature engineering for ML approaches. We believe that physical models can aggregate agronomic knowledge that ML algorithms eventually cannot capture yet directly from raw

data. Finally, as different physical models can potentially capture different aspects of reality, we recommend using multiple physical models, even for similar tasks. Then ML will hold the task of capturing valuable information from the features provided by the physical models to deliver more precise water need estimation.

INTEGRATING DIFFERENT STAKEHOLDERS

The approach taken is the result of a joint and largely interdisciplinary effort, and the authors' ambition is to generate an impact on a wide community of stakeholders with the envisioned innovation. On one hand, the interoperability at the communication level provided by IoT is a key factor to promote a platform culture among not-computer-scientists, as it brings easy node deployment, data collection, and inter-researcher interaction. On the other hand, only with our flexible architecture approach due to the heterogeneities embedded in the agricultural scenario — heterogeneity of devices and simulation tools, but also farms and the stakeholders themselves — can the barriers be properly handled. Different stakeholders speak their own languages: for example, soil-moisture sensor data are numbers for computer scientists, bits for telecommunication professionals, voltage signals for electronic engineers, while the end users expect volumetric soil moisture values. As the calibration of these sensors is soil-type-dependent, geologists, agronomists, and other researchers must take part in the game of this context-dependent calibration process.

Altogether, the smooth interplay between actors with different skills and habits is a key success factor. Each stakeholder needs specific and mostly mobile services, with the appropriate human-machine interface, if we want to deploy the appropriate level of automation with the man in the loop, as required in today's agriculture. Services need to be organized like a chain of tools that mutually exchange information and understand each other thanks to a shared information model based on emerging ontologies. This shared data model fosters smooth and sustainable innovation because the tool chain may easily be extended to provide new capabilities and value propositions, and attract new stakeholders.

IMPACTS

FOR IOT PROFESSIONALS

Our approach incorporates the ML pipeline into the IoT continuum by using a structure of services deployed as containers that exchange messages through the FIWARE NGSI unified data model. This scheme impacts IoT platform development and deployment in many aspects:

- Automation: The platform provides a subscribe/notify mechanism for building automated data pipelines.
- Traceability: The storing of meta-data information about the model specifications and context of data used in the estimation allows keeping track of model forecasts, as well as quality indicators.
- Pluggability: The integration of new or updated models is facilitated by the unified data model, consuming data and producing water need estimates in a standard way.
- Flexibility: The pluggability allows IoT professionals to compose different data workflows flexibly by using various components.
- Hybrid environments: The architecture allows the use of different ML frameworks to train models, as well as hardcoded physical models.

All these aspects have been allowing a relatively simple deployment of our platform in all SWAMP pilots, each one with its characteristics and different specific goals. In Italy, the goal is to use farm data for water management and distribution (i.e., to share data outside the farm to create an even bigger system-of-systems). In Spain, the goal is to explore the limits of flexibility and precision in irrigation by going into a very fine-

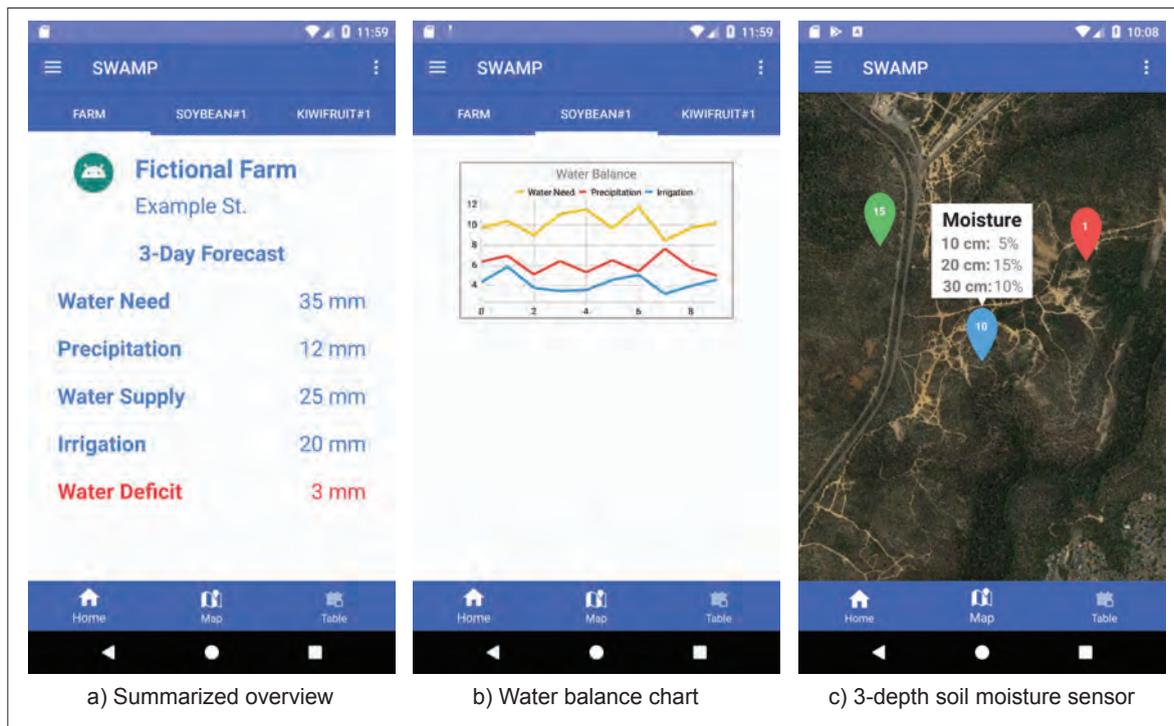


FIGURE 4. Farmer application screenshots.

grained irrigation system where each sprinkler is an IoT node. In Bahia, Brazil, there is a large-scale use case with huge center-pivot irrigation systems, where the goal is to decrease operational costs through improved situational awareness. Near São Paulo, Brazil, the goal is to improve the quality of grapes and wine.

FOR FARMERS

Currently, in modern farms that rely on physical water need estimation models and respect soil variability, farmers are provided with irrigation plans for long periods, such as weeks, months, or even the entire season. Based on their accumulated experience and daily work in the field, they continuously adapt the irrigation plan to avoid crops suffering from water stress. In this scenario, the irrigation plan plays the role of an offline longer-term forecast that needs to be fitted into the reality of the farm. In a fully automated future agriculture scenario, the IoT-based smart irrigation system will precisely control every aspect of the use of water, adapting it to shorter-term periods according to instantaneous information coming from the field, which can be daily or even based on intra-day micro adjustments in the irrigation plan.

Between these two scenarios — current and future — lies a new IoT-enabled reality that will change the way farmers face irrigation. In any case, a requirement is that farmers always control the irrigation and are provided a wealth of real-time information to be able to make better decisions. To this end, we developed a smartphone app (Fig. 4) where farmers are informed of immediate water needs (Fig. 4a), measured and forecast water balance time series (Figure 4b), and the current soil moisture information for a 3-depth sensor probe (Fig. 4c). With this real-time status of the farm at hand, and equipped with the optimized irrigation plans computed by the system, the farmer can achieve better use of the water resources without harming productivity.

As the system reliability and precision increase and earn the trust of farmers, they can slowly give more power to the system to make automated decisions. In other words, the application will allow farmers to express policies on how to behave whenever a new irrigation plan is generated.

CONCLUSION

In this article, we present our flexible IoT-ML platform and highlight its scientific contribution over related work. The platform allows easy solution deployment involving IoT and ML components working in an application. Our real case is a smart irrigation application, where we exemplify how a solution can be built and customized depending on site-specific needs. Special attention was given to how the platform enables more exploration of ML-based solutions and on how it can positively impact IoT professionals' and farmers' needs. SWAMP project pilots have just been deployed; they are operating properly, and data is being collected. The next step, expected by the end of 2020, is to analyze the data and disseminate quantitative impact results.

ACKNOWLEDGMENTS

This research has been jointly funded by the European Commission in Europe and MCTIC/RNP in Brazil, under the EUB-02-2017 IoT Pilots call.

REFERENCES

- [1] W.-L. Chen *et al.*, "AgriTalk: IoT for Precision Soil Farming of Turmeric Cultivation," *IEEE Internet of Things J.*, vol. 6, no. 3, June 2019, pp. 5209–23.
- [2] M. Roopaei, P. Rad, and K.-K. R. Choo, "Cloud of Things in Smart Agriculture: Intelligent Irrigation Monitoring by Thermal Imaging," *IEEE Cloud Computing*, vol. 4, no. 1, Jan.-Feb. 2017, pp. 10–15.
- [3] LoRa Alliance Technical Committee, "LoRaWAN 1.1 Specification," Oct. 2017; <https://lora-alliance.org/resource-hub/lorawan-specification-v11>, accessed 13 Oct., 2019.
- [4] OASIS, "MQTT Version 5.0. OASIS Committee Specification," 2 May 2018; docs.oasis-open.org/mqtt/mqtt/v5.0/cs02/mqtt-v5.0-cs02.pdf, accessed 13 Oct., 2019.
- [5] FIWARE, "FIWARE Open Source Platform," 2019; www.fiware.org, accessed 13 Oct., 2019.
- [6] L. Roffia *et al.*, "Dynamic Linked Data: A SPARQL Event Processing Architecture," *Future Int.*, vol. 10, no. 36, 20 Apr. 2018.
- [7] J. W. Jones *et al.*, "Brief History of Agricultural Systems Modeling," *Agricultural Systems*, vol. 155, July 2017, pp. 240–54.
- [8] M. Bittelli, G. S. Campbell, and F. Tomei, *Soil Physics with Python*, Oxford Univ. Press, 2015.
- [9] G. Antolini *et al.*, "CRITERIA Technical Manual," Jan. 2016; www.arpae.it/cms3/documenti/_cerca_doc/meteo/software/criteria_2016_technical_manual_english.pdf, accessed 13 Oct., 2019.
- [10] A. Nugroho, S. Hartati, and K. Mustofa, "Vector Autoregression (VAR) Model for Rainfall Forest and Isohyet Mapping in Semarang-Central Java-Indonesia," *Int'l J. Advanced Computer Science and Applications*, vol. 5, no. 11,

2014.

- [11] F. Kratzert *et al.*, "Rainfall-Runoff Modelling Using Long Short-term Memory (LSTM) Networks," *Hidrology and Earth System Sciences*, vol. 22, May 2018, pp. 6005–22.
- [12] C. Brouwer, K. Prins, and H. M., "Irrigation Water Management: Irrigation Scheduling," 1989; www.fao.org/tempref/agl/AGLW/fwm/Manual4.pdf, accessed 13 Oct., 2019.
- [13] A. Fazlali and M. Shourian, "A Demand Management Based Crop and Irrigation Planning Using the Simulation-Optimization Approach," *Water Resources Management*, vol. 32, no. 1, Jan. 2018, p. 67–81.
- [14] W. Liu *et al.*, "A Survey of Deep Neural Network, Architectures and Their Applications," *Neurocomputing*, vol. 234, Apr. 2017, pp. 11–26.
- [15] C. Kamienski *et al.*, "Smart Water Management Platform: IoT-Based Precision Irrigation for Agriculture," *Sensors*, vol. 19, no. 2, Jan. 2019, p. 276.

BIOGRAPHIES



Rodrigo Togneri (rodrigo.togneri@ufabc.edu.br) is a professor of artificial intelligence and big data in the São Paulo School of Business Administration (EAESP) of the Getúlio Vargas Foundation (FGV), and is also a D.Sc. candidate in information engineering at the Federal University of ABC (UFABC). He was a senior advisor in data science at Dell EMC, where he led several big data analytics deployment projects. His current research interests include the Internet of Things, big data analytics, and forecasting.



Carlos Kamienski (cak@ufabc.edu.br) is a full professor of computer science at UFABC, where currently he also holds the position of head of the NUVEM Strategic Research Group. He received a Ph.D. degree in computer science from the Federal University of Pernambuco (UFPE) in 2003. His current research interests include the Internet of Things, smart agriculture, smart cities, network softwarization, and fog computing.



Ramide Dantas (rasd@cin.ufpe.br) is a professor at the Federal Institute of Technology of Pernambuco (IFPE). He received his M.Sc. and Ph.D. degrees in computer networks from the Informatics Center (CIn) of UFPE in 2005 and 2012, respectively. His current research interests include the Internet of Things and machine learning.



Ronaldo C. Prati (ronaldo.prati@ufabc.edu.br) received his M.Sc. and Ph.D. degrees in computer science from the University of São Paulo in 2003 and 2006, respectively. He currently is an associate professor in the Center of Computer Science, Mathematics and Cognition at UFABC. He was a visiting researcher at the University of Bristol and the University of Granada in 2004 and 2017, respectively. His research interests include machine learning and data science.



WATER4CROPS, FIT4REUSE, GREEN4WATER).

Attilio Toscano (attilio.toscano@unibo.it) is a hydraulic engineer with a Ph.D. in environmental engineering, and an associate professor of agricultural hydraulics and watershed protection at the University of Bologna. His research activity is mainly focused on sustainable agricultural water management, precision irrigation, and NBSs (constructed wetlands) for wastewater treatment and wastewater reuse in agriculture. He was/is involved in national and European projects as WP leader or coordinator (SWAMP, MEDIWAT, CHEM-FREE,



Juha-Pekka Soininen (juha-pekka.soininen@vtt.fi) received a D.Sc. (Technology) degree from the University of Oulu in 2004. He is currently a principal scientist at VTT Technical Research Centre of Finland and a coordinator of the SWAMP project. His current research interests include the Internet of Things, cyber-physical systems, and situational awareness.



Tullio Salmon Cinotti (tullio.salmoncinotti@unibo.it) graduated in electrical engineering from the University of Bologna in 1974. He is an associate professor with the School of Engineering and Architecture at the University of Bologna, where he oversees courses on computer architecture and logic design, and is a member of the board of the Ph.D. program in Structural and Environmental Health Monitoring and Management. His research interests range from embedded systems to semantics-based data distribution architectures for systems of systems.

FOOTNOTES

- ¹ SWAMP (www.swamp-project.org) is a European Union and Brazil research partnership that holds a pilot project to explore alternatives for implementing IoT into different types of environments and irrigation systems [15].
- ² All websites cited throughout the text have been checked at Oct. 14 2019.
- ³ Physical models process mainly weather data to estimate water consumption, and as consequence, the amount of water to be replaced.
- ⁴ Soil moisture sensors are monitored. When the water content approaches a critical level, an automated trigger starts the irrigation process.
- ⁵ Evapotranspiration is the main water consumption physical process, combining evaporation from soil and plant transpiration [7].
- ⁶ Panel VAR (Vector Autorregressive) and RNN-LSTM (Recurrent Neural Network, using Long Short-Term Memory architecture) are time series machine learning techniques.

PRECISION AQUACULTURE

Fearghal O'Donncha and Jon Grant

ABSTRACT

Precision aquaculture is founded on a set of disparate, interconnected sensors deployed within the marine environment to monitor, analyze, interpret, and provide decision support for farm operations. Recent technological innovations facilitate aquaculture becoming part of the Internet of Things (IoT) — modern farms are characterized by hundreds of interconnected sensors that store and serve data, interact with other sensors and devices, and connect with a fog and cloud ecosystem. We describe the implementation of the precision aquaculture concept to a number of farms in eastern Canada. The work combines partners from industry, technology, and academia to provide data-driven insight and decision that promotes ecologically sustainable intensification of aquaculture. The article presents a first case study on how IoT can instrument, inform, and impact the aquaculture industry. Challenges related to connectivity, interoperability, and standardization are discussed, and we elucidate how our experiences can inform future activities.

INTRODUCTION

Aquaculture, or the farmed production of fish and shellfish, has grown rapidly, from supplying just 7 percent of fish for human consumption in 1974 to more than half in 2016. This rapid expansion has led to challenges including concerns over environmental degradation, disease and parasite outbreaks, and the need to efficiently manage resources to maximize productivity. These factors are pushing farms toward more efficient management practices aimed at the sustainable intensification of the industry. At the same time, innovative technologies are making the collection, processing, and analysis of large volumes of heterogeneous datasets possible. Taken together, these two factors are empowering a precision aquaculture framework that combines sensors, cloud, and analytics to enable real-time, evidence-based decision making to optimize operations.

Precision aquaculture [1] involves a variety of sensors used to gain insight into the farm environment, make decisions that optimize fish health, growth, and economic return, and reduce risk to the environment. This trend parallels developments in agriculture, where sensors and other observing technologies lead to enhanced insight into crop health as well as animal welfare. The fundamental approach has been summarized as a series of steps, namely observe, interpret, decide, and act [1]. Traditionally, many of these steps have required human intervention and depended heavily on farmer experience and intuition for correct decision and action. As farm size increases, however, and moves further offshore, automation is imperative to enable economically feasible operations.

Materialization of precision aquaculture depends on IoT technologies to empower management in a chaotic environment subject to the vagaries of oceans and weather. An obvious impediment is water cover, but other major obstacles exist, including the harsh environment, power and connectivity in offshore locations, large range of spatial scales involved (fish-, cage-, farm-, and bay-scale), and the challenges of manual intervention or analysis in the ocean (where access can be regularly impeded or prevented by adverse weather). A fish farm has an imposing array of underwater chains, ropes, moorings, and other infrastructure, so wireless communications are essential. Further, the distributed nature of the industry, composed of a large number of small-scale aquaculture companies and sensor providers, poses challenges related to the integration of diverse, sometimes proprietary, datasets into a unified edge, fog, and cloud ecosystem.

Application of mature monitoring, modeling, prediction, and analysis tools to aquaculture farms has potential to improve operations and alleviate key challenges facing the industry. Fish feed represents 50–70 percent of fish farmers' production costs, while the growth rate of fish is intrinsically linked to feed composition and time of supply; precise management can

link fish growth with optimal feed schedule and composition that minimize waste (and subsequent pollution of surrounding waters) and improves productivity. Disease and parasite-induced impacts are a major issue for aquaculture farms, costing the industry up to \$10 billion annually and having severe socio-economic impacts. Further parasite control treatment in salmon farms constitute 7.5 percent of total production costs [2]. Farming in the open ocean requires the ability to respond to natural fluctuations that impact operations, such as dissolved oxygen (DO) concentrations and temperatures, both of which act as stressors, impact feeding and parasitic rates, and even cause mortalities. Today, management of most of these tasks is conducted manually, relying on direct human observation or human-centric data acquisition means to observe conditions, combined with decision making based on subjective experience. However, as real-time sensor technologies become more prevalent on farms, the foundation exists to transition the industry from ad hoc decision making based on heuristics and intuition to real-time informed decisions backed by artificial intelligence (AI) insights and IoT connectivity.

This article describes a precision fish farming framework we have implemented on farms in Canada, which is also rapidly being implemented in Europe. It is part of an ongoing effort to develop a prototype, open-standards-based ecosystem that combines monitoring, modeling, insight, and decision making toward an autonomous framework to manage farms. It represents a multi-disciplinary collaboration with partners from the aquaculture industry, academia, and technology.

WHAT IS PRECISION AQUACULTURE?

The rapid development of aquaculture in recent years has been likened to a "Blue Revolution" [3] that matches the "Grain Revolution" of higher cereal yields from the 1950s onward. The industry's rapid growth and expansion globally, however, has caused concerns about negative environmental impacts, such as eutrophication of nearby waters and habitat alteration. In Europe, annual growth of aquaculture has declined to 1 percent, partly because of market factors, but also because the industry is subject to stringent regulation regarding sustainable development. These factors have led to a strong focus on the ecological development of aquaculture in marine systems, and the promotion of terms such as "ecological aquaculture" and "eco-aquaculture." Coupled with the need for greater efficiencies and economies of scale to empower the sustainable growth of the industry, precision aquaculture focuses on exploiting modern technologies toward the eco-intensification of aquaculture farms.

Data generated on modern aquaculture farms extend across a wide variety of forms. In situ sensors sample large numbers of environmental variables such as temperature, current velocity, dissolved oxygen (DO), chlorophyll, and salinity. Remotely sensed environmental data can sample much larger spatial domains and can be at the bay scale — from land-based sensors

Digital Object Identifier: 10.1109/IOTM.0001.1900033

such as CODAR-type HF radar — or at the global scale from a satellite-based monitoring system. Informing on farm operations also requires sampling of animal variables such as size, clustering behavior, and movement, and this is typically done using underwater technologies such as video monitoring, hydroacoustic technology, and aerial drone imagery.

Further, there are large datasets of pertinent variables that are generated by numerical models such as weather or ocean circulation products. These datasets constitute huge data volumes with distinct characteristics. Integrating and extracting information from these disparate data sources are key to encapsulating the full dynamics of the farm environment and enabling effective management. Related data from mathematical models are estimates of fish growth and behavior that can be used to guide expected conditions and decision [4].

The overarching aims of precision aquaculture have been defined as [1]: 1) improve accuracy, precision and repeatability in farming operations; 2) facilitate more autonomous and continuous biomass/animal monitoring; 3) provide more reliable decision support; and 4) reduce dependencies on manual labor and subjective assessments, thus improving staff safety. Similar to precision livestock farming [5], precision fish farming has been decomposed into three conditions that must be fulfilled. We note that in addition to these caveats, we include sensing of the ambient environment (e.g., water temperature, oxygen), a consideration that is less important in agriculture where animals can be housed. The basic requirements of precision aquaculture are:

- Continuous monitoring of animal variables (i.e., parameters related to the behavioral or physiological state of the fish)
- A reliable model to predict how animal variables dynamically vary in response to external factors
- Observations and predictions integrated into an online system for decision or control

Achieving these objectives is dependent on the successful implementation of a range of innovative technologies related to sensors, computer vision, and AI, enabled by a readily interconnected edge, fog, and cloud ecosystem. Central to this paradigm shift from human to autonomous management is an IoT platform to link information from different components, understand current status against a desired or model-predicted benchmark, and return insight from data in terms of actionable information, such as modified feeding protocol or defined health intervention or treatment.

Conceptually, the cultivation of fish in the ocean has parallels with terrestrial livestock farming. In practice, however, livestock farming is more amenable toward direct human and animal interaction than is possible in the marine-based counterpart. Modern fish farms comprise cages with up to 200,000 fish. As farms are typically composed of 10–20 cages, and multiple farms are often co-located in a bay, the total number of individual fish is enormous. This precludes the direct translation of concepts from livestock farming, and in practice, precision aquaculture is a marriage of approaches developed for both precision livestock and grain cultivation; that is, fish are not managed as individuals as are cows, but are obviously more complex in management than plants.

DEEPSense FOR AQUACULTURE

DeepSense (<http://www.deepsense.ca>) is a big ocean data innovation environment, powered by IBM, that brings together academia and industry to drive growth in the ocean economy. A key component is the commercialization of IoT technologies toward better management of fish farms. Specifically, a new research program involving Dalhousie University, DeepSense, InnovaSea, Cooke Aquaculture, and IBM has been created to research sensor networks, big data, and analytics applied to fish farming in eastern Canada. Dalhousie University in Halifax, Nova Scotia is a global leader in the marine sciences and aquaculture, and home of DeepSense in the Faculty of Computer



FIGURE 1. Sensor configuration within a stylized cage. Nine sensors were deployed within each cage (we only show five to illustrate approximate locations) consisting of four sensors at 2 m depth in each corner (north, south, east, and west), four sensors at 8 m depth in each corner, and one sensor at 4 m depth in the center of the cage. This sensor density is for research purposes of understanding spatial variation in the net pen. Operational metrics related to fish position and behavior are estimated using the “CageEye” acoustic system as well as a number of individually tagged fish.

Science. The university collaborates with InnovaSea, also headquartered in Nova Scotia. Cooke Aquaculture is an international seafood company originating in New Brunswick, Canada, with a deep commitment to innovation and sustainability, cooperating closely in research with Dalhousie. The unique combination of industry, technology, and scientific expertise further positions Nova Scotia as a global center of ocean technology, developing innovative solutions to empower aquaculture operations.

FARM MONITORING

Within this precision fish farming initiative, hundreds of real-time underwater wireless acoustic sensors have been deployed in Canada at multiple fish farms by Cooke and InnovaSea (<http://www.rtaqua.com>). Sensors take 100,000 measurements daily, analyzing 11 million data points about temperature and tilt, salinity, dissolved oxygen, blue-green algae, chlorophyll, and turbidity. Figure 1 presents a schematic of the sensor deployment that collects pertinent environmental variables within a cage. Additional data on fish position are provided by the “CageEye” acoustic system (<http://www.cageeye.no>), as well as individually acoustically tagged fish.

All data generated on farms are communicated to IBM® Cloud (<https://www.ibm.com/cloud>), utilizing the open-standard Message Queuing Telemetry Transport (MQTT) protocol for data transport. For each cage, a comprehensive set of variables are collected, communicated, and updated, continuously informing on environmental and animal conditions.

The ocean consists of complex environmental conditions (tides, winds, water masses, ice) that impact farm operations, safety, and health of the fish. Hence, information external to the cage is pertinent to operations and management. Satellite measured observations, weather data, and numerical models of the ocean all generate information impacting at the farm scale. Real-time analysis and decision making require the ability to rapidly query and extract pertinent variables from these datasets. We integrate in situ and geospatial datasets using a big data platform, Physical Analytics Integrated Repository and Services (PAIRS) [6], a service that processes petabytes of data and

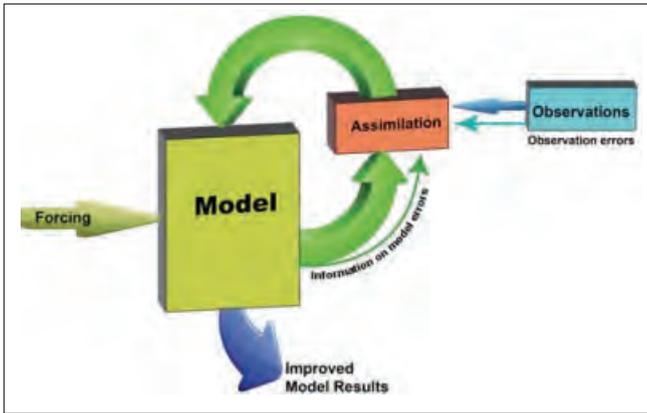


FIGURE 2. Mechanistic models contain errors that increase with time due to model imperfections and deviations of forcing conditions from reality. Data assimilation minimizes these errors by correcting the model stats using new observations (from [10]).

addresses the spatial and temporal complexity associated with heterogeneous data integration. Built on top of the open source big data technologies Hadoop and HBase, PAIRS aims to accelerate data queries by curating and storing geospatial datasets from diverse sources (NOAA, NASA, ECMWF, etc.) in a scalable storage table that can be rapidly accessed and retrieved.

MODELING WITHIN PRECISION AQUACULTURE

The objective of precision aquaculture is to manage the observed status of the farms relative to a defined benchmark (e.g., projected biomass). Hence, a key functionality of the IoT platform is the capability to manage various machine learning models and integrate with the different data streams coming from sensors, weather data, and other open sources. A range of machine learning and mechanistic models relate to managing aquaculture operations. In particular, we focus on:

- Mechanistic and data-driven models to predict fish health, biomass, and mortality based on information on feed and environmental stressors
- Predictive models to inform on outbreaks of parasitic infections
- Deep learning models to forecast oceanographic conditions multiple days/weeks in advance

Fish feed is the most expensive part of aquaculture and causes environmental problems when excess product sinks to the bottom. Optimal supply of feed is a complex selection that includes feed composition, growth stage, biomarkers, and environmental conditions. While mechanistic models have been developed that simulate growth rates based on feeding regime and environmental conditions, the nonlinear relationships and sensitivity to external events such as diseases or parasites have made prediction difficult [7].

A practical solution requires that prediction be based on observed status to maintain accuracy. One approach combines mechanistic models with observations using data assimilation, a mathematical technique that incorporates process knowledge encapsulated in a physics-based model with information from observations describing the current state of the system (described schematically in Fig. 2). A precision aquaculture implementation can be summarized as:

- Dynamic process models for individual fish growth based on feeding regime and environmental conditions are implemented.
- Continuous update of model state based on actual fish position and/or biomass as measured by the CageEye system described earlier or the Biosonics aquaculture biomass monitor (<https://www.biosonicsinc.com/products/aquaculture-biomass-monitor/>) is performed.

Within the former approach, data assimilation concepts have seen enormous application since the 1960s as scientists aimed to update models using sparse sensor observations [8]. As sensors become more prevalent, data-intensive computing is continuing to transform industries and decision making [9]. Leveraging the large datasets being generated on aquaculture farms has multiple advantages, particularly related to extracting insight from highly complex nonlinear processes not amenable to encoding within a set of explanatory equations. An obvious case in aquaculture is fish health and in particular parasitic outbreaks.

Sea lice presence in salmon farms is a complex interplay of hydrodynamics, lice load, temperature, and position of the fish in the water column. Nonlinear, opaque relationships have traditionally made mechanistic modeling impractical. More recently, IBM, in collaboration with industry stakeholders, has implemented a deep learning model that collates data from multiple sources and predicts sea-lice outbreaks, termed “AquaCloud.” The model was fed with data on environmental conditions and lice counts from over 2000 salmon cages along the Norwegian coast. Combining a dense network of environmental sensors and manual sampling (of lice count), the deep learning model provides two-week-ahead prediction of lice count with 70 percent accuracy [11]. Within a precision aquaculture framework, advance prediction of parasitic outbreaks presents opportunities for improved management and treatment that can reduce severity of outbreak and invasiveness of treatment.

Machine-learning-based models for geophysical processes are an active area of research. The authors recently developed and demonstrated a machine learning surrogate model for a physics-based ocean-wave model [12]. The machine learning model yielded enormous speedup (> 5000-fold) in computational time while maintaining accuracy that was well within the confidence bounds of the physics-based model. In effect, deep-learning-based approaches enable the transition of complex modeling systems from HPC to edge devices (naturally, the training of the models is expensive, but once trained, deployment is cheap). This approach is being extended as part of DeepSense with data from hundreds of sensors being fed to deep learning models that provide continuous prediction of oceanographic variables multiple days in advance. A number of studies have produced promising results using machine learning to predict pertinent variables such as ocean temperature [13] and algal blooms [14].

FROM DATA TO DECISION

The key objective of precision aquaculture is moving beyond data toward decision. As part of the DeepSense platform, an IoT network has been developed to integrate data from hundreds of sensors at salmon cages in Canada. This is complemented by a model management framework that enables tracking of models and functionalities, automatic subscription to data streams, and relationships between different models (geospatial, vertical dependencies, etc.). Efforts are ongoing to integrate this with an evidence-based decision platform. Currently, we focus on two key challenges facing fish farms:

- Optimizing feeding to maximize productivity and minimize environmental impacts
- Inform on health intervention practices to mitigate sea lice

Optimizing fish feed needs to consider the composition and schedule in response to external conditions. The objective is the supply of nutritionally appropriate feed at a rate and frequency that maximizes uptake by the fish. Some guidelines may instruct — such as not to feed when environmental stressors may impact consumption — but ultimately real-time conditions and behavior need to inform the decision.

The rate of supply of feed can be related to the monitoring of cage biomass and activity. Namely, when monitoring activity indicates that feeding behavior has concluded (i.e., the fish

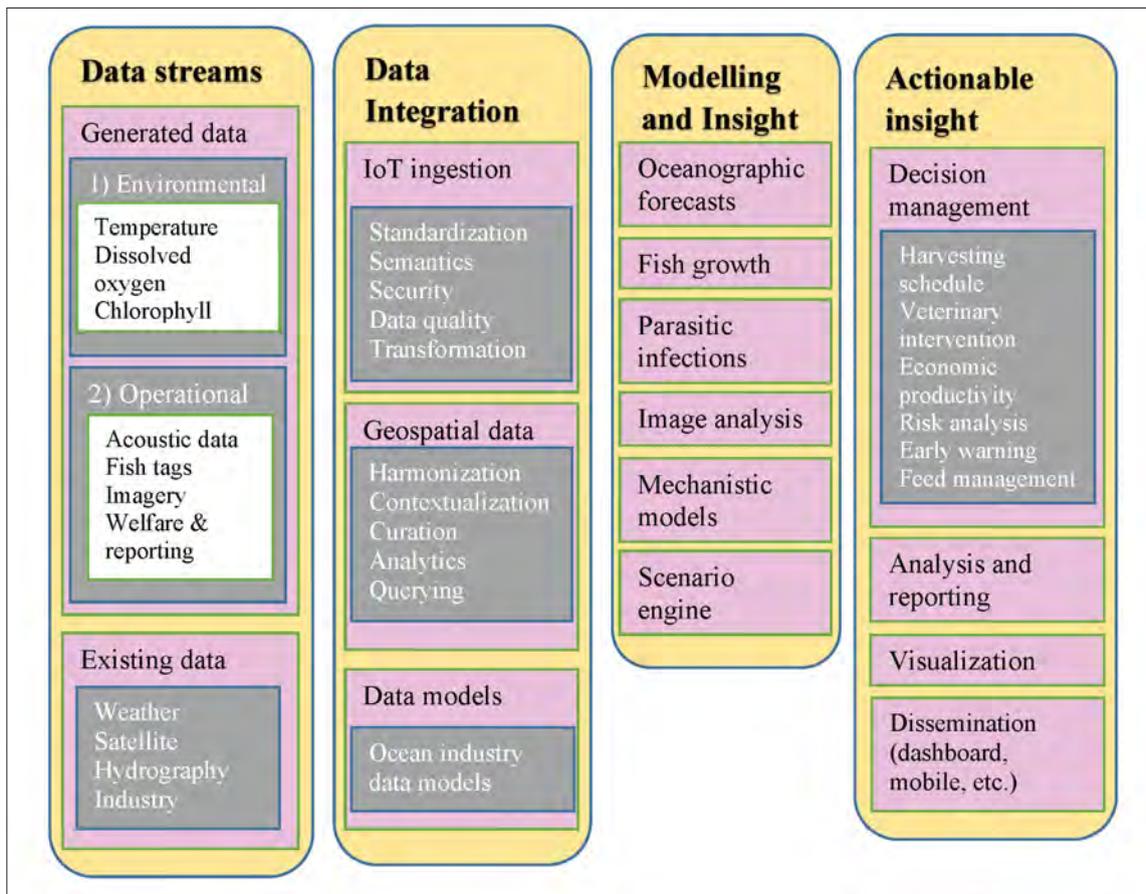


FIGURE 3. Schematic of the precision aquaculture framework encompassing the in situ monitoring of farm conditions and operations, integration of the generated and existing data in a unified cloud platform, analytics and machine learning applied to the data to generate insight, and the dissemination of that data to stakeholders in an actionable format.

move away from the surface where feed is supplied), feeding is stopped to prevent waste and environmental pollution. This is implemented via an AI system that processes information from video or the acoustic monitoring sensor to label observed data as “feeding” or “not feeding.”

As with terrestrial farming, there is extensive knowledge on the most appropriate feed composition at different stages of fish growth, health, and seasonal cycle. A key challenge is applying this knowledge to real-world scenarios in the face of uncertainty. An IoT solution provides continuous update of measured fish size against predicted values, enabling response to deviations. Namely, an online AI system monitors current biomass and recommends the most appropriate feed composition based on a database of growth conditions and nutrient requirements.

Parasite and pathogen outbreaks have traumatic impact on farms, leading to mass mortalities, causing fish to be unmarketable, and generating huge damage to public perception. As is often the case, prevention is preferable to a costly cure that is dependent on harsh chemicals or highly invasive mechanical removal. Predictive models, as described above, allow for advanced treatment to avoid these symptoms, reducing fish loss and economic losses.

Figure 3 presents a schematic of the different components of a precision aquaculture framework. It can be broadly decomposed into four pillars: the monitoring of environmental and operational conditions at the cage, farm, bay, and ocean scales (considering both in situ generated data and existing data from sources such as NASA Modis or ECMWF); integration of data from available sources into an accessible form; applying models and analytics on the data to generate insight; and the dissemination

of that insight to stakeholders in an actionable format (directly to farm operators, summary metrics to management, report generation for regulatory requirements, etc.).

THE FUTURE OF PRECISION AQUACULTURE

Fish farming is a relatively young industry but, in some ways, has been quicker to adapt to difficult circumstances than land-based farming because of modern technology. The next phase of industrialization is dependent on using data to inform decisions. Certain challenges exist related to its location in the ocean – requiring robust, low-cost sensors capable of underwater and in-air wireless connectivity. However, the industry has seen huge progress in this regard with many farms being equipped with a dense network of sensors streaming data in real time. Similar to other industries, the current focus is extracting actionable insight from IoT data [15].

Interoperability poses a significant challenge as sensors currently cover a wide range of types, suppliers, and levels of sophistication. This extends from legacy sensors storing data in onboard data loggers to modern sensor stacks reporting in proprietary format to dedicated cloud platforms. DeepSense is committed to an open standards approach based on MQTT protocol. Extensive work is ongoing with sensor manufacturers as well as the aquaculture industry more broadly to standardize messaging protocols. These include activities we are developing as part of Horizon 2020 project GAIN (Green Aquaculture Intensification; https://www.unive.it/gainh2020_eu) and previous work conducted by IBM with seven different Norwegian aquaculture companies as part of the AquaCloud project for sea lice data. Security and sovereignty of data is critical to fully

exploit AI capabilities in an ethical and commercially sustainable way. An often overlooked part is the content returning from sensors and empowering analytics to understand and process these messages. Interoperability of these messages with agnostic IoT platforms requires insight into what the value or content of each sensor refers to via semantic domain models, for example [16].

Computer vision is currently receiving a lot of attention at the academic and venture capital level. Proponents claim that computer vision and AI can be used to monitor feeding behavior and fish biomass, detect sea lice, and optimize the supply of feed, medicine, and other resources to farms [17].

At its core, precision aquaculture is dependent on leveraging IoT technologies to move beyond data toward insight. By integrating data from heterogeneous, disparate sources into a unified cloud platform, it promises to move from heuristics and experience toward evidence and information. Aquaculture is projected to supply 62 percent of fish for human consumption by 2030, and securing this supply is contingent on eco-intensification of the industry based on data.

CONCLUSION

Because precision fish farming is in its early stages, the development and proliferation of sensors is a growth area. A wide variety of sensors are feasible, including optical, acoustic, and biological sensors for currents, particles, pathogens, and harmful algal blooms. Moreover, a similarly diverse array of image-based data are being applied to fish farming ranging from direct videography of fish to satellite remote sensing. The use of drones in data capture is an obvious application of airborne technology. The attendant development of AI to analyze images and interpret essential information related to fish behavior and health is an active area of research. While the benefits of these advances in husbandry are apparent, their application to public facing indicators of sustainability is critical. The expansion of big data in fish farming should have spinoffs for a larger conversation regarding indicators of sustainability in aquaculture.

ACKNOWLEDGMENTS

The study was funded by a grant to JG as the NSERC-Cooke Industrial Research Chair in Sustainable Aquaculture at Dalhousie University and the European Union's Horizon 2020 programme as part of the RIA GAIN project under grant agreement No. 773330. We thank our colleagues in GAIN and at Cooke Aquaculture, Evangelos Milos of DeepSense, Brad Rogers and Tim Stone of InnoSea, Cal Gosse of IBM, and Ramon Filgueira, Meredith Burke, and Caitlin Stockwell of Dalhousie University.

REFERENCES

- [1] M. Føre et al., "Precision Fish Farming: A New Framework to Improve Production in Aquaculture," *Biosys. Eng.*, vol. 173, Sept. 2018, pp. 176–93.
- [2] A. P. Shinn et al., "Economic Costs of Protistan and Metazoan Parasites to Global Mariculture," *Parasitology*, vol. 142, no. 1, Jan. 2015, pp. 196–270.

- [3] B. A. Costa-Pierce, *Ecological Aquaculture?: The Evolution of the Blue Revolution*, Wiley, 2008.
- [4] M. Føre et al., "Precision Fish Farming: A New Framework to Improve Production in Aquaculture," *Biosys. Eng.*, vol. 173, Sept. 2018, pp. 176–93.
- [5] D. Berckmans, "Automatic On-Line Monitoring of Animals by Precision Livestock Farming," *Livestock Production and Society*, R. Geers and F. Madec, Eds. Wageningen Academic Publishers, 2006, pp. 27–30.
- [6] L. J. Klein et al., "PAIRS: A Scalable Geo-Spatial Data Analytics Platform," *IEEE Int'l. Conf. Big Data*, 2015, pp. 1290–98.
- [7] M. Føre et al., "Modelling Growth Performance and Feeding Behaviour of Atlantic Salmon (*Salmo salar* L.) in Commercial-Size Aquaculture Net Pens: Model Details and Validation through Full-Scale Experiments," *Aquaculture*, vol. 464, Nov. 2016, pp. 268–78.
- [8] M. J. Martin et al., "Status and Future of Data Assimilation in Operational Oceanography," *J. Oper. Oceanogr.*, vol. 8, no. sup1, Apr. 2015, pp. s28–s48.
- [9] A. Hey, S. Tansley, and K. Tolle, *The Fourth Paradigm: Data-Intensive Scientific Discovery*, Microsoft Research, 2009.
- [10] P. Houser, "Improved Disaster Management Using Data Assimilation," *Approaches to Disaster Management-Examining the Implications of Hazards, Emergencies and Disasters*, IntechOpen, 2013.
- [11] T. Hoel, "Data Science Helps Norway's Fish Farmers Keep Salmon Populations Healthy," <https://www.ibm.com/blogs/cloud-computing/2018/09/17/data-science-norway-fish-farmers/>, accessed 14 May 2019.
- [12] S. C. James, Y. Zhang, and F. O'Donncha, "A Machine Learning Framework to Forecast Wave Conditions," *Coast. Eng.*, vol. 137, 2018.
- [13] S. Wolff, F. O'Donncha, and B. Chen, "Statistical and Machine Learning Ensemble Modelling to Forecast Sea Surface Temperature," Sept. 2019.
- [14] F. O'Donncha and S. C. James, "A Machine-Learning-Based Approach to Forecast Harmful Algal Blooms," *Aquaculture Europe*, 2019.
- [15] "IDC European Vertical Markets Survey," 2017.
- [16] J. Ploennigs, J. Cohn, and A. Stanford-Clark, "The Future of IoT," *IEEE Internet of Things Mag.*, vol. 1, no. 1, 2018, pp. 28–33.
- [17] M. J. Coren, "Aquaculture Is the Latest Hotness for Silicon Valley Venture Capitalists in 2018 — Quartz," <https://qz.com/1256208/aquaculture-is-the-latest-hotness-for-silicon-valley-venture-capitalists-in-2018/>, accessed 14 May 2019.

BIOGRAPHIES



Fearghal O'Donncha is a research scientist at IBM Research — Ireland. His work focuses on applying simulation models, analytics, and machine learning techniques to assist industry operations. This encompasses developing and deploying simulation-based models, integrating sensor data from a variety of IoT platforms, developing AI-based models that extract value from sensor or expert data, and optimizing these tools to a variety of HPC and cloud-based platforms. He is an adjunct faculty member at the National University of Ireland, Galway.



Jon Grant is a professor of oceanography at Dalhousie University. He is the NSERC-Cooke Industrial Research Chair in Sustainable Aquaculture, in a multi-year partnership with Cooke Aquaculture, one of the major global aquaculture and fisheries companies. Trained as a benthic ecologist, he has a B.Sc. from Duke University and a Ph.D. from the University of South Carolina. He has worked in aquaculture-environment interactions for 30 years and is co-editor of a new book, *Goods and Services of Marine Bivalves*.



Join the Technology Communities Around the World at IEEE Conferences!

Are you interested in learning the latest technologies and meeting global communications experts? Annually IEEE Communications Society (ComSoc) holds over 40 conferences around the world that provide technical communities many opportunities to hear world-class expert insights, technology trends and advancements, and network with leaders of various fields of expertise. Conferences presentations cover topics including 5G, AI, ML, wireless communications, sensor and space technologies, IoT and more.

Check out the following conferences and join today!

IEEE ISPLC 2020

IEEE International Symposium on Power Line Communications and its Applications
31 March-3 April 2020 // Malaga, Spain
<https://isplc2020.ieee-isplc.org/>

IEEE WCNC 2020

IEEE Wireless Communications and Networking Conference
6-9 April 2020 // Seoul, South Korea
<https://wcnc2020.ieee-wcnc.org/>

IEEE WF-IoT 2020

IEEE 6th World Forum on Internet of Things
5-9 April 2020 // New Orleans, Louisiana, USA
<https://wfiot2020.iot.ieee.org/>

IEEE NOMS 2020

IEEE/IFIP Network Operations and Management Symposium
20-24 April 2020 // Budapest, Hungary
<https://noms2020.ieee-noms.org/>

IEEE INFOCOM 2020

IEEE International Conference on Computer Communications
27-30 April 2020 // Beijing, China
<https://infocom2020.ieee-infocom.org/>

IEEE ICBC 2020

IEEE International Conference on Blockchain and Cryptocurrency
3-6 May 2020 // Toronto, Canada
<https://icbc2020.ieee-icbc.org/>

IEEE HPSR 2020

IEEE International Conference on High Performance Switching and Routing
11-14 May 2020 // Newark, New Jersey, USA
<https://hpsr2020.ieee-hpsr.org/>

IEEE CTW 2020

IEEE Communications Theory Workshop
17-20 May 2020 // Banff, Canada
<https://ctw2020.ieee-ctw.org/>

IEEE ICC 2020

IEEE International Conference on Communications
7-11 June 2020 // Dublin, Ireland
<https://icc2020.ieee-icc.org/>

IEEE CNS 2020

IEEE Conference on Communications and Network Security
29 June-1 July 2020 // Avignon, France
<https://cns2020.ieee-cns.org/>

IEEE NetSoft 2020

IEEE Conference on Network Softwarization
29 June-3 July 2020 // Ghent, Belgium
<https://netsoft2020.ieee-netsoft.org/>

IEEE PIMRC 2020

IEEE International Symposium on Personal, Indoor and Mobile Radio Communications
31 August-3 September 2020 // London, UK
<https://pimrc2020.ieee-pimrc.org/>

IEEE GLOBECOM 2020

IEEE Global Communications Conference
7-11 December 2020 // Taipei City, Taiwan
<https://globecom2020.ieee-globecom.org/>



IEEE

IEEE ComSoc
IEEE Communications Society

www.comsoc.org

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 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].

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

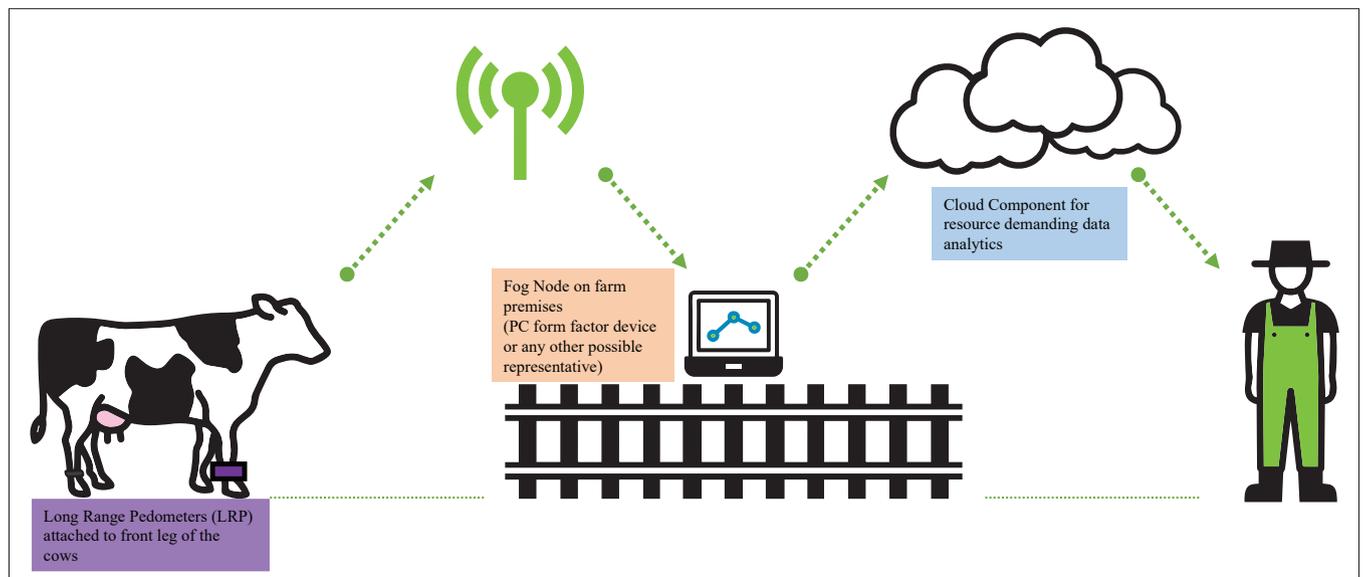


FIGURE 2. System workflow and different components of the developed IoT solution.

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 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.

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,

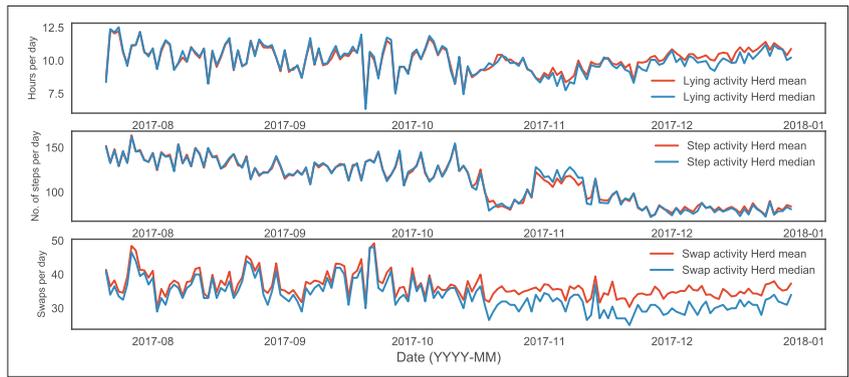


FIGURE 3. Comparing the mean and median of the various animal activities.

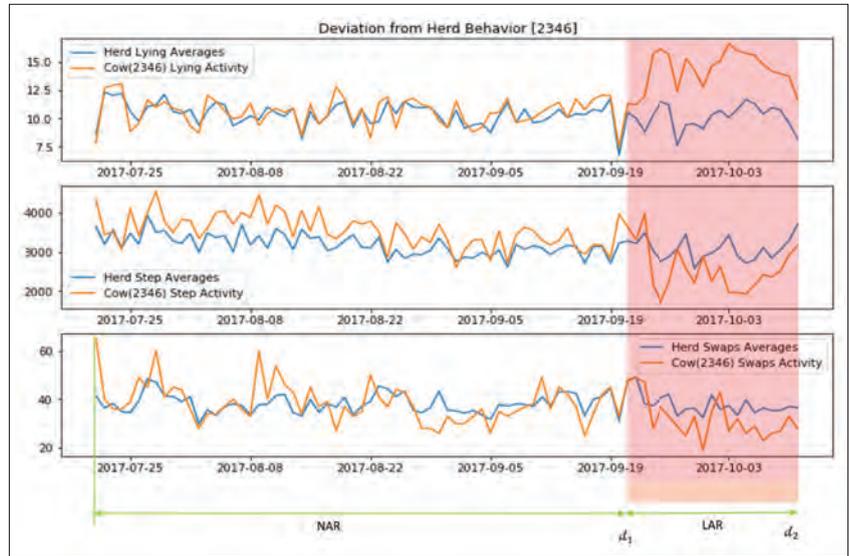


FIGURE 4. Relationship between herd mean and cow activity for cow 2346, showing deviation in its behavior from the herd as it transitions into lameness.

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 **Machine Learning Based Early Lameness Detection 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/#.WdTbDmhSyUI>, accessed Dec. 14, 2017.
- [3] H. Kharas, "The Emerging Middle Class in Developing Countries," OECD Development Centre Working Papers, 2010, <http://dx.doi.org/10.1787/5k-mmp8lncrns-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. Golasowski, and D. Timmerman, "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.*, "Smarterd-Connected Cows demo.mp4 – Google Drive," <https://drive.google.com/file/d/1Q1rKSp8SkAZRRAFQDQHdPXu1T-VYyV3/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 Emerging 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 Telecommunications, 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.

GROWING PLANTS, RAISING ANIMALS, AND FEEDING COMMUNITIES THROUGH CONNECTED AGRICULTURE: AN IOT CHALLENGE

Ankita Raturi and Dennis Buckmaster

ABSTRACT

The Internet of Things is a growing field of design and development in agriculture. In this article, we provide IoT researchers and practitioners a glimpse into the motivations, needs, and challenges faced when designing digital technologies for agriculture. We describe three farming scenarios and offer a vision for the power of IoT in agriculture, followed by a discussion of opportunities for design. We build the argument for why just collecting data isn't enough and suggest target areas for the design of ubiquitous digital technologies for agriculture. Finally, we introduce four communities of practice on IoT for agriculture.

INTRODUCTION

Agriculture is core to human survival. Through agriculture we harness the power of nature to grow plants and raise animals to feed, clothe, and fuel our communities. From axe to drone, we have designed and developed technologies to augment our physical and mental faculties in the pursuit of food, and eventually, the practice of cultivation and husbandry. The design of agricultural technologies is not new, but our methodologies, farming practices, communities and values, science, and environment are all changing. Agricultural and technological practice and innovation are deeply intertwined.

The dawn of agriculture is marked by our transition from nomadic to settlement-based societies in the Neolithic period, as we began to domesticate plants and animals. Since then, there have been three technological innovations that have catapulted our agricultural capacity. First was the industrial revolution, spawned by the development of the John Deere steel plow that magnified our ability to break the soil and enable improved seedbeds at that time. Subsequent development of tractors and other machines transformed agricultural operations, making agricultural mechanization one of our top 10 feats of engineering [1]. In the early 1900s came chemical innovations exemplified by the discovery of the Haber-Bosch process: an artificial nitrogen fixation technique used to produce ammonia. Our ability to produce synthetic fertilizers dramatically increased food production. Genetic engineering marks the third wave of technology in agriculture: Norman Borlaug's development of wheat varieties for multiple growing seasons in a year, disease resistance, and, most famously, dwarfism. Along with the continued progress in cultivation methods, these have resulted in the massive expansion of food production throughout southeast Asia, popularly known as the Green Revolution. Each of these periods of technological innovation have had complex ripple effects on human nutrition, environmental conditions, economic prosperity, social justice, and the very nature of agriculture.

We are now well into a new era in agriculture that builds on the knowledge and tools from each of those earlier revolutions: the digital revolution. We design technologies to help us manage complex food logistics and distribute food fairly; improve worker livelihoods, community well being, and equitable money distribution; improve nutritional quality of food, while enabling creativity in cooking; improve animal welfare, soil health, and air and water quality, and create more regenerative agricultural systems. Agricultural technologies are no small matter of programming, and the challenge that lies ahead for the IoT commu-

nity is a transdisciplinary problem in connecting plants, animals, machines, people, and environments, to support resilience in our food system.

This has inevitably led to the design of a plethora of networked devices for sensing and actuation, conceptualized as the Internet of Things (IoT), with the promise of technologies to empower us to sense more, act swiftly, and make better decisions. Early visions of IoT for agriculture focused on the augmentation of human senses. Cameras and imaging provide sight at a distance, with computer vision enabling the detection of crop health and tracking of animal movement. Chemical traces can be sniffed out using air quality monitors, allowing for the monitoring of methane emissions by animals in confined environments. Probes allow us to touch soil to sense for moisture, and our hearing is augmented through devices that listen for the presence of predators in our pastures. However, every sensor we introduce runs the risk of inducing sensory overload. An increased interest in sensing farms introduces both opportunities for new agricultural practices but also a cognitive overload on farmers, consumers, and everyone in between, as they are faced with a glut of data. The near-term challenge in agricultural IoT is to consider: How can we empower agricultural stakeholders with high quality and timely data for better decision making?

At minimum, IoT is simply a network of sensors and actuators deployed in a given context [2]. IoT involves machine-to-machine interaction, where each machine may consist of a data acquisition component, a computational and data storage component with networking capabilities, and sometimes, actuators or control logic. Taken to its logical extreme, IoT involves ubiquitous computing, with sensors and actuators embedded in our landscapes, abstracting and automating certain categories of actions (e.g., moving things) and decisions (e.g., when to turn a switch on) without human intervention. It is imperative, therefore, that we consider community values, civil liberties, the future of work, and environmental impacts, among other consequences when designing IoT systems, including IoT for agriculture. Technology is an alloyed good, which combines potential benefits to society and human well being with externalities that may only appear post-deployment.

In the coming decade, the challenge will be to consider: How can we harness the power of digital agricultural technologies to improve, sustain, and grow with care? Innovations in this space have applications throughout our food system, from agriculture, to production, transportation, processing, marketing, consumption, all the way to waste management. The thoughtful implementation of IoT in agriculture offers the radical opportunity to improve resilience in our food system and enable data-driven regenerative agriculture.

Digital Object Identifier: 10.1109/IOTM.0001.1900105

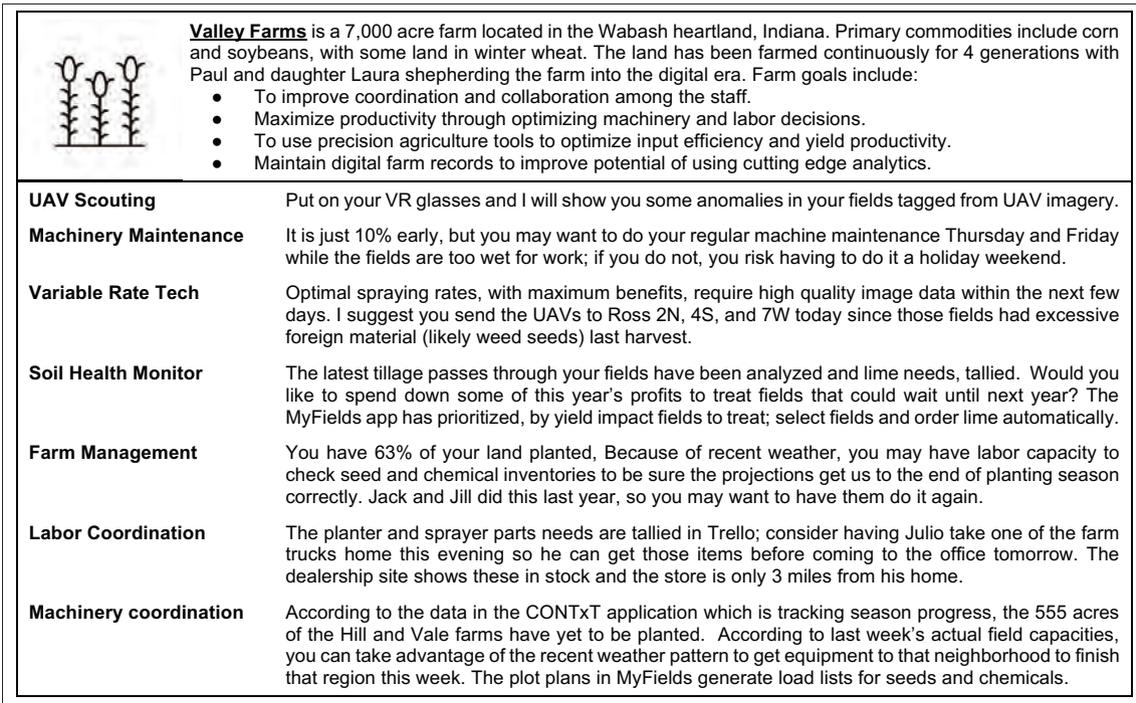


FIGURE 1. Valley Farms: A design scenario envisioning IoT in agriculture.

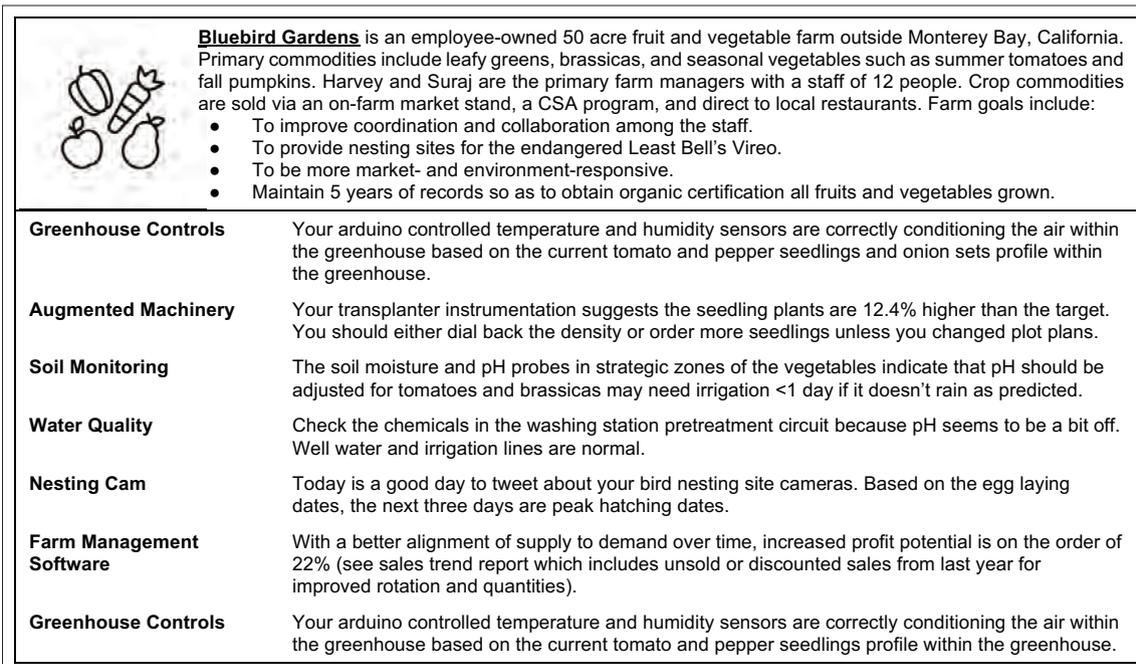


FIGURE 2. Bluebird Gardens: A design scenario envisioning IoT in agriculture.

A TALE OF THREE FARMS

In each of Figs. 1, 2, and 3, we detail a design fiction scenario to offer a sampling of agricultural complexity and decision making. Each scenario is concerned with an archetypal farm and describes a suite of IoT sensing and intelligent decision support that explore the untapped but reachable potential of IoT in agriculture.

CONNECTED AGRICULTURE

We describe five categories of connectedness in agriculture: plants, animals, machines, people, and environments. By breaking down the landscape of agricultural IoT in this manner, we offer a variety of entry points for IoT researchers and

practitioners to consider when designing sensors, actuators, micro-computers, and IoT devices for agricultural use cases.

CONNECTED PLANTS

We grow a very wide diversity of plants as agricultural commodities. Each type of plant has its own unique management challenges, structural properties, tolerances, and components that we care about the most. The efficacy of plant sensing technologies varies widely. Highly standardized plants grown in monocultures offer the least amount of variability and have proven to be a good testbed for many plant sensing technologies. As sensors and our ability to use such data with robotics improves, we are beginning to see more and more design for specialty crops (vegetables, fruits), including interest in design

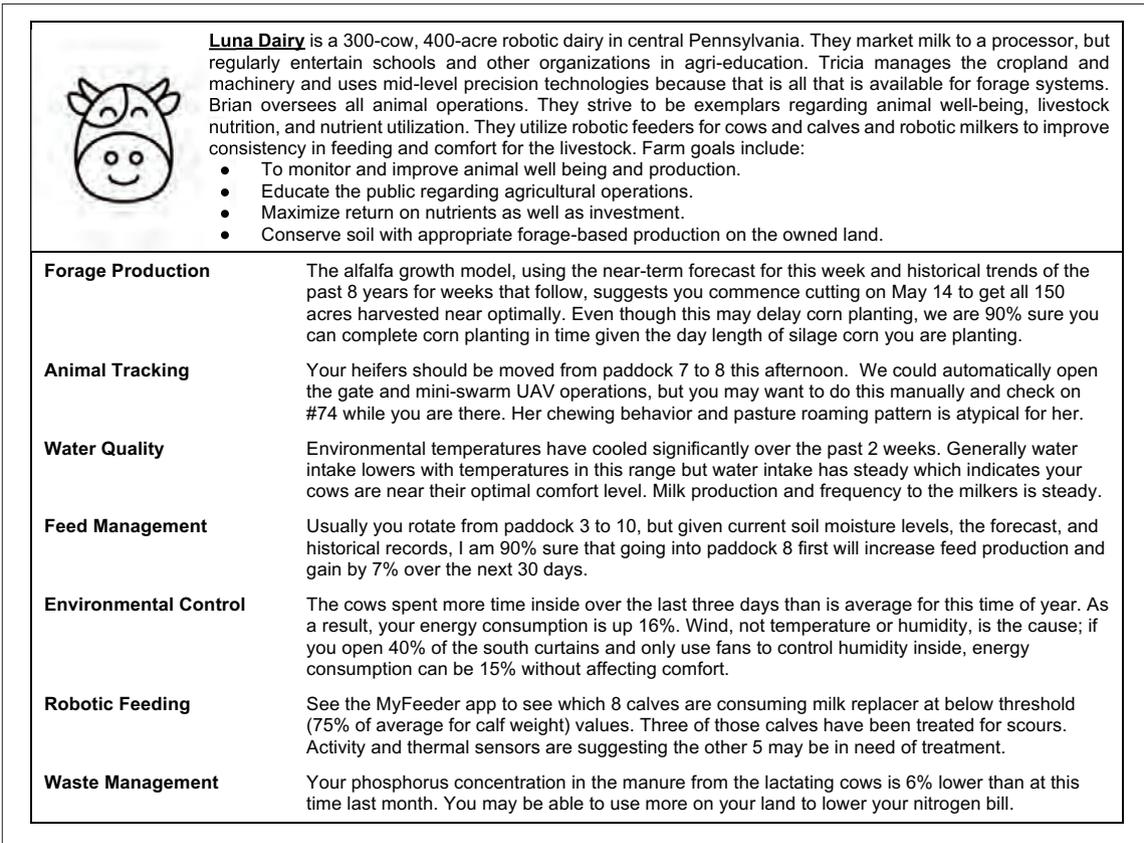


FIGURE 3. Luna Dairy: A design scenario envisioning IoT in agriculture.

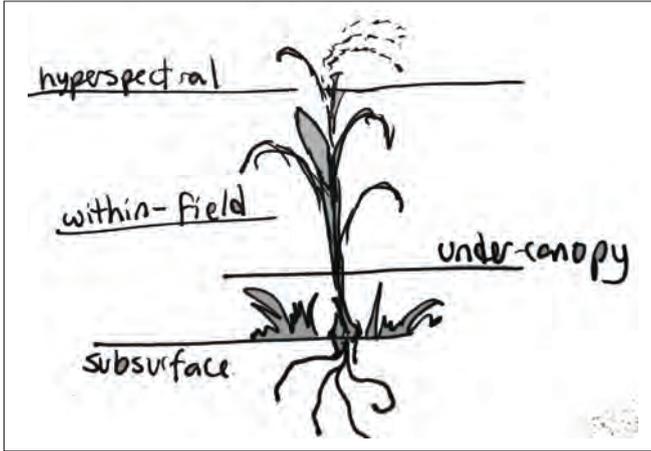


FIGURE 4. Plant sensing

for the management of diverse farming systems.

Figure 4 overviews the different “levels” at which we sense plants. Each of these types of sensor data, when coupled with an actuator, may be used for a variety of strategic and tactical decisions in a field. Some examples include:

- Subsurface: Sensing root architecture can inform our assumption about nutrient uptake, and thus determine nutrient management and tillage remediation.
- Under-canopy: By counting plant stems in a row, we can determine yield potential and gain an understanding of how plant density affects plant growth or planter settings affect germination.
- Plant-level: Genomics data and within plant sensing via, for instance, hyperspectral imaging and high throughput phenotyping offer insights into the interplay between genetics and environment.

- Overhead: Machine or unmanned aerial vehicle (UAV) mounted sensors can be used to determine weed-crop competition or disease manifestation, providing farmers with a map of crop health to inform field management.
- Extraterrestrial: Remote sensing can be used reasonably effectively to estimate biomass or ground cover and subsequently crop effects on the landscape.

CONNECTED ANIMALS

Livestock introduce many levels of complexity when thinking about IoT for agriculture. While raising animals for food, there is a delicate balance between managing their diet, environment, and other factors to result in high-quality meat. There is also a need to care for their well being and quality of life. Animal needs vary widely; to date, IoT for livestock has focused on support for beef and dairy cows.

- Individual (external). A common practice, whether animals are kept indoors or outdoors, in water or on land, is tagging. Tags have evolved from brands to clips on ears, to RFID tags that allowed for tracking of animal movement, and more recently are beginning to include a variety of sensors. For example, accelerometers placed on ear tags can provide researchers with valuable data on animal behavior, allowing early detection of stress, sickness, estrus, or pregnancy.
- Individual (internal). More recently, experimental technologies are being developed to provide livestock managers with insights into individual animals’ internal health as well. This includes, for instance, the placement of small devices inside the rumen of a cow to monitor their digestive activity to determine animal health and feed efficiency and to anticipate effects on the resulting meat or milk.
- Groups. Herd tracking is also a growing area of interest, with explorations in the use of drones to track, guide, and potentially deliver medicines to animals that roam in pastures. The idea behind these systems is to minimize human

intervention in the life of the animals, conceivably reducing stress and allowing for increased independence of grazing herds.

CONNECTED MACHINES

Given early visions of IoT as simply machine-to-machine interaction, it makes sense that agricultural machines were some of the first things to be networked. Early digital agriculture took the form of precision agriculture, which primarily involved automation of farm activities typically conducted by machinery. The introduction of drive-by-wire meant that tractors, harvesters, combines, and other large farm vehicles were the perfect testbed for early IoT (with much of their data flowing in standard controller area networks, CANs [3]). Variable rate technologies exemplify early success in closing the sense-decide-act loop in agriculture [4]. For example, variable rate planting technologies can utilize a soil map (generated with remotely sensed images and/or machine data that is augmented with ground truth data collected by field scouts) to plant different crop varieties in different soil types. Subsequently, a variable rate applicator could, using a projected yield map based on historical data, apply fertilizer in a more optimal manner. While such technologies have been relatively widely adopted, variable rate technologies and other precision machinery have typically been used in the realm of fairly large-scale, wide-acre monoculture cropping systems.

There is also growing interest in creating IoT systems for indoor livestock agriculture, where animals are able to interact with a variety of machinery from weighing scales, feeding stations, to milking devices. Here, dairy cow operations have received the most attention with dramatic impacts on the landscape of dairy, particularly in the United States. An instrumental approach has been robotic milking and robotic feeding.

CONNECTED ENVIRONMENTS

Weather, soil, air, and water quality monitoring, for both indoor and outdoor agriculture, are increasingly important as our environmental conditions become more volatile. Some of the first forms of sensing to be streamed from farms include rainfall, humidity, temperature, and other data from weather stations located on farms across the world. At times, these are farmer-owned and operated, while in other cases a weather station may be part of, for instance, the National Weather Service network. Land and water on farms are typically sampled periodically, sent to laboratories where they are tested, for example, for microbial, chemical, and nutrient composition. More recently, soil and water sensors have been developed to, for example, detect soil moisture, pH, and nutrients. Such sensors are increasingly connected to networked computers that provide access to real-time environmental data at specific locations on farms. Such data are crucial components of agricultural IoT as they provide critical site conditions that often determine the constraints and conditions for various agricultural activities to be conducted (e.g., smart irrigation).

A systems approach to connected agriculture is currently best exemplified by vertical agriculture. Crops are grown in carefully controlled greenhouses, with nutrients delivered through a network of pipes and filters, controlled lighting and air conditions, and constant streams of data about plant and growth chamber properties available to farm managers via a suite of dashboards.

CONNECTED PEOPLE

Most farmers and farm workers in the United States are already connected via smartphones. Currently, there are many applications to allow people to track their work, but also coordinate and collaborate on farm activities. Some applications provide real-time streaming data collected by sensors (e.g., weather station apps), while others provide real-time location of tagged livestock. In many ways, the current state of IoT means that for each type of thing connected in a farm, there is likely a

standalone web or mobile application that people must use to interact with the data. However, there is significant untapped potential for wearable technologies that are enabled with voice input and smart algorithms to provide hands-free, and ideally automatic, data collection, manipulation, and visualization of agricultural data.

INFRASTRUCTURAL LIMITS

It is critical to note that there are two fundamental hardware limitations to current efforts in IoT for agriculture. First are power limitations. The lifetime of a battery is particularly important as the frequency of change must be considered. In the case of large-scale agriculture, IoT sensors may be deployed across vast spaces. As the distances between plants, animals, machines, and people are great, the frequency of battery change as a result of battery life is particularly problematic. A farmer does not want to chase down a cow to replace its ear-tracker battery or have to visit each sensor; this would introduce an entire layer of maintenance due to IoT device density. Since plugging in an IoT device is not always an option in agriculture, many devices are designed to include, for instance, small photovoltaic systems to produce their own power. However, limitations on panel size and efficiencies constrain the computational capacity of such devices.

Second is Internet connectivity in rural spaces. Agricultural landscapes are less sparsely populated than urban areas. For too long, this has been the rationale used to excuse limited to no availability of broadband Internet in rural communities [5]. The promise of IoT in agriculture to provide decision support based on real-time site-specific conditions is hampered by our ability to transmit data on farms: from big data (e.g., drone-captured imaging of vast grazing lands), to distributed data (e.g., location and other data from hives located on orchards across a region), and dense data (e.g., multiple, by minute measurements of water quality in aquaponics farms). While there is growing support for increased access to rural broadband, we argue that edge computing, mesh networking, and continued reduction in cost and size of micro-computers can still allow for IoT innovation in agriculture. Furthermore, new technologies and approaches for connectivity, sometimes with delay or low bandwidth, solve some of these problems.

OPPORTUNITIES FOR DESIGN

In the last decade, there have been some key technological developments that improve our ability to realize IoT for improved resilience in food systems and enable data-driven regenerative agriculture. There has been a steep decline in the cost of sensors, micro-computers, actuators, and other IoT components, along with a growing interest in developing environmental sensing technologies for smart and connected cities, ecologies, and agriculture. In turn, this has led to research efforts in data science, machine learning, ontologies, and decision support to take advantage of increased availability of agricultural data. An increasing interest in open source technologies, as well as demand for access and agency to one's data, further has the potential for innovation in agricultural IoT. Opportunities and challenges for design are driven by urgent issues faced in agriculture due to climate change, support for rural communities, increasing inequality, and public values and interest in the provenance of food.

The reality and future of IoT in agriculture will unfold soon. We offer eight opportunities in design for improved research and innovation in IoT for agriculture.

Design for Sustainability: We have previously issued a call to action, bringing together through human-computer interaction (HCI) researchers, designers, and practitioners to critically engage in the design of technologies for a more sustainable food system [6]. We argue that human-centered, community-oriented, and environmentally sensitive approaches to research and development are critical for IoT adoption in agriculture.

Design for Rural Communities: The increasing presence of connected devices in our homes and workplaces will also be matched with increasing familiarity and comfort with such technologies in general, which may influence adoption of agricultural IoT. In the case of agriculture, the digital divide remains within countries including the United States, where there is a gap in access, availability, and education regarding digital technologies between urban and rural communities [7]. We have previously described infrastructural challenges in agricultural landscapes. Digital transformation must also be paired with access to employment and education opportunities for workforce development and modernization. If the future of agriculture is digital, we must prepare for the shift in skills required for the future of work.

Design for Data Sovereignty: Pursuant to achieving its technical potential, IoT can be thought of as ubiquitous computing, with sensors integrated throughout our landscapes. However, monitoring without agency or control is simply surveillance. Digital agricultural technologies must consider the data rights of farmers while offering consumers and other farm stakeholders appropriate insights into the provenance of food. To achieve this, we must negotiate the balance between transparency and accountability with privacy. This is further complicated as networked landscapes inevitably introduce security threats into otherwise isolated systems. We argue that successful IoT for agriculture must begin with transparency in data use, a focus on the ethical implications of monitoring and control, and consideration of data sovereignty of food system stakeholders to be a primary requirement.

Design for Agricultural Diversity: Agricultural systems are extremely diverse, as are the factors of inspiration and drive toward IoT implementation for different crops, animals, practices, and places, where initially, IoT systems were not commonly available for small-scale agricultural systems; nor were they particularly affordable. However, the increasing proliferation of low-cost sensors and microcontrollers, matched by consumer appetites for understanding the provenance of their food, has led to an explosion of interest in digital technologies across systems and scales [8].

Design for Trust and Accountability: We have previously called for the design technology to increase trust and accountability in our food system [6]. IoT, in particular, offers the capacity to ground-truth agricultural practices, by allowing for provenance data to truly begin within the farm itself. Indeed, there are several current efforts in research and practice proposing the harmonization of farm sensor and/or sample data, and farm management information, in the context of environmental conditions to provide verifiable evidence of sustainability claims. Indeed, IoT in agriculture is a foundational component of improving traceability in the supply chain, particularly given the current trend of utilizing blockchain technologies for supply chain verification.

Design for Agricultural Practice and Decision Support: Effective IoT for agriculture demands the closing of the sense-decide-act loop. Achieving dramatic improvements (true loop closing) will require interoperability unlike ever seen before in agricultural contexts. Data from machinery, sensors in soil, products, bins, sensors on animals, data regarding workers, weather, imagery, audio traces, and video need to first be interoperable and machine readable; then true fusion can occur. Researchers wrangle data (and do it in post-processing mode) in order to test hypotheses and develop models, but practitioners certainly do not have the time or generally the skill set to gather, reformat, align, visualize, and analyze data.

Given our silo-ized sensor-driven approach to IoT, to date we have typically only been able to offer decision support for the optimization or management of single variables (e.g., yield optimization, irrigation timing, weather reporting). Early developments in one-dimensional data-driven approaches in agriculture were technologies for soil characterization, crop health, yield

maps, and livestock feed management. Some popular “full-stack” technologies include variable rate applicators for more efficient input use (e.g., water, pesticides, fertilizer) in large-scale row cropping systems and, more recently, robotic milkers and feeders that allow many cows and calves to find their own milking rhythm and schedule. The next level of decision support must have larger reach.

In addition to the often one-dimensional approach, current commercial IoT and analytic platforms largely focus on strategic decisions. A couple of exceptions would be controls for irrigation, grain bin aeration/drying, and greenhouse conditions. Because of this, there has been little need for real-time data flow (or maybe it is that the lack of real-time connectivity has resulted in the delay-tolerant focus). With a dramatic improvement in interoperability and connectivity, there is great potential to also improve tactical decisions. The current state of delayed access to data for insights, however, often results in poor or missing data [9]. Such poor and occasionally sparse data leaves lots of potential for improving decision making. If we can achieve connectedness and interoperability, the state of the firm/farm/operations can be known, so managers can better make the “best next” decision. In cropping systems, this might entail knowing moisture content of soils in all fields and stage of growth of crops in all fields. In livestock systems, this might be knowing daily rate of gain and daily feed intake per animal. In logistics, it could be up-to-date knowledge of queues and at-the-moment capacities.

We look forward to a near-term day when economic and environmental sustainability (and also logistics when it comes down to details) and optimal nutrient management in cropping systems can be accomplished using relevant data. That may include public data (topography, soil type), non-IoT data (soil sample results), and sensor data to improve related decisions. By integrating georeferenced yield data (i.e., “fused” analysis from sensors for location, speed, width, crop flow, crop moisture), electron conductivity of soil, and UAV imagery, which can indicate plant health and nutrient concentrations, we will be able to improve strategic decisions within the context of an entire season. We will also influence tactical decisions and implementation over time with people and machines, through optimization (even by simplifying constraints or using simplified objective functions).

Design for Ubiquity: As we are beginning to think about sensor consolidation to offer multi-sensed models of farms and critically consider the diversity of commodities in agricultural systems, we are able to design for agricultural diversity. Effective IoT in agriculture requires intelligently deployed sensing, which also captures structured metadata (i.e., full context) in machine readable formats with the potential for human comprehension.

If agricultural data streams and sets really met the FAIR (findable, accessible, interoperable, and reusable) principles (even in a private context), tactical decisions could be facilitated with microservices and apps that merge data and models. The three farming scenarios depicted earlier in this article hint at some of these possibilities. Consider, for example, that if soil type, soil cover, weather, and topography were known, and soil moisture and growing degree days (GDD) could be computed, visualized, and analyzed. This could positively influence decisions regarding sequencing of spring work in fields; if planting date and variety (GDD to maturity) were automatically recorded, GDD tracking could provide an approximate status of each field going into each next “phase” of the growing season (scouting, spraying, harvesting). Combined with aerial imagery or sensor data that kept these models on target, a farmer would have decent assurance of near optimal logistics.

Design for Interoperability: Digital agriculture is in its infancy and currently involves a fragmented landscape of data, models, tools, and communities. Several efforts exist to introduce conceptual and practical interoperability to enable seamless IoT-based systems. For instance, the Open Technology Ecosystem

for Adaptive Management [10] effort around interoperability in field-level measurement techniques aims at a connected user experience through an ecosystem of tools and a data sharing community for the enablement of soil health research.

Such efforts require ontologies, application programming interface (API) frameworks, and standards as fundamental building blocks for interoperability for use in communities of practice that include farmers, researchers, and developers alike. For example, the Open Ag Data Alliance [11] is an open source extension of the REST framework designed for agricultural data interoperability. OADA provides a “standard API framework for automated data exchange” with an immediate focus on knitting together disparate machine data streams.

The harmonization of sensor data requires an ontological consistency among farm models and software. We have previously developed the Modeling Sustainable Systems (MoSS) framework for modeling complex adaptive systems for modeling farms [12]. The goal was to provide an information model of an agricultural system that used a coherent vocabulary and syntax, mapped onto farmers’ mental farm models, and enabled the spatio-temporal representation of heterogeneous farm data to enable design for decision support. Our current work includes the extension of MoSS as a means to harmonize sensor data across scales.

OPEN SOURCE FOR IOT IN AGRICULTURE

An open source approach to software development democratizes innovation, removes barriers to collaboration, increases markets, and improves the talent pipeline [13]. A functioning and interoperable middle layer — between the raw sensor data and insight to users, or better yet automated controls — can only be achieved via open source development where standards naturally emerge due to success in achieving the goal [2]. This same open source culture speeds innovation because there is less reinventing of interfaces, conversion utilities, and algorithms. It results in more productive development due to a talent pool knowledgeable in how to FIND solutions to seemingly new problems from other fields.

We invite the IoT community to four communities of practice on agricultural technology. We introduce these communities simply as an entry point for IoT researchers. As we are founding members of the groups listed in this section, we describe this selection as we can offer a point of entry for IoT researchers and practitioners looking to engage in IoT development for agricultural use cases.

The Gathering for Open Agricultural Technologies (GOAT): A grassroots, online, open source community, the GOAT forum and instant messaging channels offer an easy place to begin: <http://forum.goatech.org>. GOAT was initially founded to bring together farmers, researchers, and technologists interested in coordinating open source development of digital technologies for agriculture, including IoT [14].

The Open Agricultural Technology & Systems Center (OATS): Researchers at the Purdue OATS are focused around a suite of topics including sensor development and harmonization, machine automation, data interoperability, human centered design, and agricultural modeling and simulation, detailed at <http://oatscenter.org>. OATS faculty argue that IoT for agriculture requires coordination across each of these fronts, across government, industry, and research through an open source development paradigm.

Precision Sustainable Agriculture (PSA): For IoT researchers particularly interested in technology targeted at large-scale sustainable agriculture research and practice, including farming techniques such as cover cropping, visit <http://www.precision-sustainableag.org>. Researchers are particularly interested in the integration of sensing techniques across scale via the consolidation of sensor data from probes, drones, and satellites.

Open Technology Ecosystem for Adaptive Management (openTEAM): A collaborative community of farms, research labs, non-profit organizations, food companies, and food system stakeholders. Members are dedicated to the development of critical technologies to improve our understanding of regenerative agricultural practices, particularly in service of adaptive soil health management. Working groups, ways to get involved, and more information can be found at <http://openteam.community>.

REFERENCES

- [1] Nat'l. Academy of Sciences, *Greatest Engineering Achievements of the Twentieth Century*, <http://www.greatachievements.org/>, accessed Sept. 12, 2019.
- [2] A. Tzounis et al., “Internet of Things in Agriculture, Recent Advances and Future Challenges,” *Biosystems Engineering*, vol. 164, 2017, pp. 31-48.
- [3] M. L. Stone et al., “R.K. ISO11783: An Electronic Communications Protocol for Agricultural Equipment,” ASAE Distinguished Lecture #23, *Ag. Equip. Tech. Conf. Publ. No. 913C1798*, American Society of Agricultural and Biological Engineers. St. Joseph, MI, 1999.
- [4] X. Pham and M. Stack, “How Analytics is Transforming Agriculture,” *Business Horizons*, vol. 61, 2018, pp. 125-33.
- [5] U.S. Dept. of Agriculture, “A Case for Rural Broadband,” 2019, <https://www.usda.gov/sites/default/files/documents/case-for-rural-broadband.pdf>, accessed Sept. 12, 2019.
- [6] J. Norton et al., “A Grand Challenge for HCI: Food+ Sustainability,” *Interactions*, vol. 24, no. 6, 2017, pp. 50-55.
- [7] N. M. Trendov et al., “Digital Technologies in Agriculture and Rural Areas — Status Report,” Rome, 157 pp., Licence: cc by-nc-sa- 3.0 igo. 2019.
- [8] Refresh Food+Tech Working Group Report, “From Soil to Supper,” 2018, <https://report.refreshfoodandtech.com/>, accessed Sept. 12, 2019.
- [9] D. Buckmaster et al., “Use Cases for Real Time Data in Agriculture,” *14th Int'l. Conf. Precision Agriculture*, Montreal, Canada, 2018.
- [10] “Open Technology Ecosystem for Adaptive Management,” <https://openteam.community/>, accessed Sept. 12, 2019.
- [11] Open Ag Data Alliance, <https://github.com/OADA/>, accessed Sept. 12, 2019.
- [12] A. Raturi, *Modeling Sustainable Agriculture*, UC Irvine, 2017.
- [13] A. Ault, J. Krogmeier, and D. Buckmaster, “Ag’s Future Belongs to Open Source,” *ASABE Resource Mag.*, Mar./Apr. 2018, pp 7-9.
- [14] A. Raturi, “The GOAT Report,” 2018, <http://goatech.org/2018/07/22/the-goat-report/>, accessed Sept. 12, 2019.

BIOGRAPHIES



Ankita Raturi is an assistant professor at Purdue University and runs the Agricultural Informatics Lab focused on community-oriented design, domain-specific modeling, and software engineering for a more resilient future. She works with GOAT, OATS, PSA, and openTEAM.



Dennis Buckmaster is a professor and Dean’s Fellow for Digital Agriculture at Purdue University. He co-founded and co-directs the OATS center, where his current research is focused around farmer-focused tactical decision making that integrates modeling, data engineering, and data science.



AROUND THE WORLD OF IOT

In this column we plan to take a tour around different physical locations in the world with the objective of highlighting the peculiarities of the trendiest IoT-related applications in selected regions. Thus, the “IoT World” will certainly be physical, but traveling around it shall also expose to the readers how different application domains have been addressed, with particular attention to business sustainability.

A DIVE INTO THE AGRITECH WORLD: TECHNOLOGIES AND ADOPTION INCENTIVES

by Raffaele Giaffreda

Chief IoT Scientist, Fondazione Bruno Kessler

In this column, rather than taking a trip to a geographical location, we explore the world of AgriTech, shedding light on the currently available technology assets and on some of the hurdles any technology transfer initiative in this domain is facing, slowing down adoption of ICT in agriculture. Through this journey we will therefore at first take a look at the landscape of technology enablers supporting the vision of an upcoming fourth agricultural revolution, while in the second part we will juxtapose a picture of what the world of potential adopters looks like, identifying what are the adoption showstoppers.

LET'S START WITH SOME NUMBERS

Precision crop farming's market size alone is set to increase from USD 1 billion in 2018 to USD 2.3 billion in 2023¹. The smart agriculture market was valued at USD 6.34 billion in 2017 and is expected to reach USD 13.50 billion by 2023, at a CAGR of 12.39 percent during the forecast period². Agriculture employs more than one billion people and generates over USD 1.3 trillion dollars worth of food annually³. AgriFood Tech startups, innovating from farm-to-fork, raised USD 10.1 billion in 2017, a 29 percent year-over-year increase⁴.

From a mere numbers perspective we are looking at a steadily growing sector. In this article we focus within the Primary Production sector of the whole food value chain, the one that should be of immediate concern for growers and farmers. We will talk about AgriTech as opposed to FoodTech: separate considerations and analysis for the food transformation and distribution sectors are outside the scope of this column.

The AgriTech is there to give opportunities for remote monitoring and precise treatments in the fields, increasing yields while going back to sustainable practices. The whole sector is on the verge of a new revolution, and that in itself is a big happening worth following. A look back is enough to realize how agricultural practices have been consistently refractory to change. In fact, from 10,000 BC when stationary farming was introduced until the present day, the rate of change in agricultural practices only had a spike between the 17th and late mid 19th centuries (new techniques for crop rotation, more efficient use of arable land and utilization of livestock) and between the 1950s and 1960s (mechanization and the use of chemicals). The above-mentioned expected growth numbers though, analyzed together with current challenges in terms of climate change, sustainability, growing population and new, low-cost technology availability, all point in the direction of a new agricultural revolution.

THE TECHNOLOGICAL ASSETS

Technology has definitely made a leap forward in the last decade toward creating solutions that are more and more affordable also for the Primary Production sector, historically characterized by low margins.

On this front innovation started already a few years ago leveraging on earth observation and remote sensing with satellites like the U.S. Landsat 8 (2013 launch)⁵ and the Sentinels from the EU/ESA program Copernicus (Sentinel-2A launched in 2015 and Sentinel-2B launched in 2017)⁶.

Copernicus in particular aims to foster innovation and business development with freely available satellite multi-spectral images, which are made available every few days. These images, duly interpreted and georeferenced, can provide a lot of useful information about the situation in a crop/soil. The granularity of this data, however, does not go to lower than 10 meters resolution and therefore AgriTech solutions leveraging such data would only be suitable for large scale farming and crops like wheat, maize, barley, rye, oilseed rape, potatoes, etc. Still a lot of information can be extracted by calculating, for example, various vegetation indices⁷. With improvements in imaging technologies on board satellites, also came the ability to map soil features, crop vigor, chlorophyll content, the need for fertilizers, water, and yield for some type of crops.

On the innovation front many startups leverage these open-data to implement and sell farm management systems that fulfil short term needs of the agriculture industry. Such data also enable unprecedented applications like JRC MARS⁸ which can support growers with regular bulletins derived from their Crop Monitoring Service. Looking further ahead, land observation satellite data also provide valuable information for policy makers to set guidelines on the proper management of this important sector (i.e., 2018-23 U.S. Agricultural Policy and 2021-27 EU Common Agricultural Policy⁹).

Remote satellite sensing can leverage open-data and certainly helps large-scale farming businesses and policy makers, but has its limitations when it comes to resolution (both in time and space dimensions), crop diversity and field size. This is where progress in “closer to earth”, yet remote sensing technologies have appeared to address some of these shortcomings.

Here the data-collection is no longer relegated to the design of an API into an open-data store, but it becomes part of the investment when one needs more precise AgriTech solutions. In recent years we saw the booming of drone-harvested data (which, like satellite, use passive remote sensing technologies¹⁰) but also the advent of more costly but highly efficient active remote sensing technologies such as LIDAR (Light Detection and Ranging) mounted on unmanned vehicles and using scanning lasers for various applications, more related to durable monitoring. Costs and complexities of these types of technologies are such that their use is more of a one-off every few years to survey the soil rather than being used in daily interactions with crops.

For highly interactive purposes, addressing the time and space lack of granularity and producing continuous data-streams, we find the IoT-based solutions: the advent of low-cost, long duration, wide coverage networks (i.e., LPWANs) for collecting data is enabling the ultimate precision agriculture solutions. These cover the holes that remote sensing cannot fulfil and provide high granularity monitoring data where the resolution needs to go well below the 10m threshold or where there is no remote visibility (such as in greenhouses and tree shaded areas etc.) and with readings well within one hour intervals (compared to “every few days” from satellite, for example).

The growth in this sector has been fuelled by reduced sizes and costs of sensing devices as well as increased coverage of data collection networks which have considerably reduced the density of fixed infrastructure needed to create meaningful services for the growers.

WHAT ARE THE BARRIERS TO ADOPTION?

Having seen a quick overview of assets and technologies available for AgriTech solutions, we use this second part of the column to review the hurdles that we still see on the path to adoption and to making the fourth revolution in agriculture happen for real.

Being on the forefront of technology transfer in this domain, we can share some of our impressions from working directly with growers. To facilitate passing the message, we draw a comparison between the Internet of People and the Internet of Things, trying to derive from the former lessons which, duly contextualized, should also apply to the latter. The objective here is to explain what barriers one has to face when innovating in the agricultural sector.

Well, much of the contextual background shows evidence of a world where farmers and growers dig their knowledge deep into their predecessors' practices. The non-millennials of our readers might still remember people resisting ideas about online shopping stores, claiming they would never give up on their "look and feel" experience while shopping. Many of us still go to food markets and visit bookstores, but the above claim hasn't stopped the likes of Amazon or Walmart from riding the success wave of online shopping, fuelled by the enthusiasm of early adopters and by reliable access and web technologies that made the whole experience stress-free.

By contrast, "it has always been done this way" is by far the answer you hear most when you question the reasons that lead to this or that agricultural practice. Or it is a justification for not being willing to bother, or being afraid of what unforeseeable problems innovation might cause. The biggest barrier to change in reality is the problem of low margins for the actors at the bottom of the value-chain of what still remains the world's largest industry¹¹. What experience are AgriTech early adopters getting? And what are the incentives that can push adoption rates to increase?

The first hurdle to adoption is what also caused the digital divide in the Internet: the problem of access. While connectivity options do exist in rural areas, having to bother with subscription charges and SIM connected devices raises the adoption threshold. Removing that hurdle with widespread LPWAN coverage, where wireless sensing devices can work "out of the box" with no cables and no installations to worry about, greatly facilitates opportunities for "try and see." Which leads me to the other problem innovators face: the Internet of People is fast-paced, and speed is part of the game for gaining users' attention on the next sleek app.

Growers and farmers cannot be cornered with arrogant apps that pretend to control their crops better than they would. Innovators must have on their plan an initial phase where the AgriTech solution is there, "just in case," but ready to provide meaningful data or pondered advice.

The widespread success of the Internet of People came with widespread access, but also with the ability to channel the wealth of data the Web was collecting via interesting web/smartphone applications perceived as adding value to the point where consumers did not mind paying for it. Similarly for AgriTech, the million dollar question becomes what can one do with data collected in the fields, that is perceived as adding value to the growers' daily activities. Readings from sensors must be calibrated to the context (soil, meteo, crop, etc.) and interpreted to become meaningful insight rather than simply left to animate fancy dashboards. Co-creation is key and the right amount of time must be allocated to this.

Getting the correct and usable solution for the growers is only one side of the incentives coin: it should help growers with remote controlling their crops' health and conditions easily, allowing them to better invest their time and physical presence in more worthy tasks.

The other side of the incentives coin relates to the outside perception of what we, as consumers, think and want from our food. On this front there certainly is a trend where sustainable agricultural practices and products can be monetized because distribution chains want it, because final customers want it and are prepared to pay additionally for it or choose it over other products which cannot make the same claims.

Even though positioned within the poorer side of the "AgriTech to FoodTech" value chain¹², Primary Production companies can leverage this particular demand to monetize higher market values for their raw products, provided AgriTech helps them show evidence about their "sustainability compliant" practices. Nothing is better than IoT technologies and widespread sensing for producing that evidence and Distributed Ledger Technologies for keeping trace and eventually redistributing value from the "close to consumers" side of the value-chain (FoodTech) back to the "close to growers" side. Examples could be sharing among the farmers savings derived from reduced electricity bills if less irrigation water is consumed or sharing part of profit increases where distribution puts on supermarket shelves products that sell more because of the certified way they have been produced¹³; but this is part of another journey.

What we say here is that the use of IoT can also be justified if it helps to more seamlessly and automatically prove adherence to various certification requirements or to a given set of production rules, which for traceability purposes is a must¹⁴. In both cases the adopter-growers can differentiate their products on the market and AgriTech is there to support them focusing on what they can do best, which is farming and not paperwork.

In conclusion, growing population, SDGs¹⁵, climate change, advances in agriculture (i.e., vertical farming and hydroponics) are certainly going to bring changes. Those operators who still believe the "it has always been done this way" mantra will keep them out of trouble are in for a shaky ride into the future with loss of competitiveness at the end of it.



RAFFAELE GIAFFREDA (rgiaffreda@fbk.eu) is a chief IoT scientist at FBK CREATE-NET, Italy. He has worked in the telecom R&D environment since the beginning of his career, focusing in the last decade on IoT and related technology transfer activities. In his role, he is now responsible for setting research and innovation directions, acquisition of funding, and the execution of a number of collaborative projects in the IoT domain. He has worked in Italy and the United Kingdom (10 years), acquiring experience in both corporate telco environments (R&D of BT and Telecom Italia) as well as in a small research organization (CREATE-NET before its merger with FBK), where the ability to acquire funding was key to ensuring continuity of operations. He is a recognized expert with a substantial record of IEEE publications and conference presentations, a patent, and various book chapters and tutorials on IoT. He is an experienced speaker and chair of IoT related events, serves as an EU reviewer, and has served on the TPCs of a number of international conferences, and he is the Editor-in-Chief of the *IEEE IoT Newsletter*.

FOOTNOTES

- ¹ <https://www.statista.com/statistics/956934/smart-farming-market-size-europe-application-vertical/>
- ² <https://www.marketsandmarkets.com/Market-Reports/smart-agriculture-market-239736790.html>
- ³ <https://www.worldwildlife.org/industries/sustainable-agriculture>
- ⁴ <https://agfunder.com/research/agfunder-agrifood-tech-investing-report-2017/>
- ⁵ <https://landsat.gsfc.nasa.gov/landsat-data-continuity-mission/>
- ⁶ http://www.esa.int/Applications/Observing_the_Earth/Copernicus/Overview3
- ⁷ <https://www.agricolus.com/en/how-to-use-vegetation-indices-from-remote-sensing-beyond-ndvi/>
- ⁸ <https://ec.europa.eu/jrc/en/mars/bulletins>
- ⁹ https://ec.europa.eu/info/food-farming-fisheries/key-policies/common-agricultural-policy/future-cap_en
- ¹⁰ https://earthobservatory.nasa.gov/features/RemoteSensing/remote_08.php
- ¹¹ <https://www.worldwildlife.org/industries/sustainable-agriculture>
- ¹² The State of European Food Tech 2019 report, from Five Seasons Ventures <https://www.foodtech.vc/>
- ¹³ https://twitter.com/sapience_eit
- ¹⁴ <https://ieeexplore.ieee.org/document/8835426>
- ¹⁵ <https://www.un.org/sustainabledevelopment/sustainable-development-goals/>

NEW SECTION FOCUSES ON SMART CITIES

First, my thanks and appreciation to Raffaele Giaffreda for pulling together this issue. Many thanks also to the authors, reviewers, and publications staff who made this issue of *IoT* possible. *IoT* is a team effort!

With this issue, we also inaugurated the Smart Cities section of *IoT*. Rebecca Hammons and Joel Myers, who were the Guest Editors for our Smart Cities issue, have been joined by Jean Rice to be the regular editors for Smart Cities. This issue's contribution, "Growing Plants, Raising Animals, and Feeding Communities Through Connected Agriculture: An IoT Challenge," by Ankita Raturi and Dennis Buckmaster, intersects the domains of smart agriculture and smart cities. The authors present a systems approach to IoT and agriculture, and describe five categories of connectedness in agriculture, ranging from plants and animals to machines, people, and environments. Many of the design challenges in smart cities are shared by smart agriculture. The authors close with the presentation of three open-source communities of practice that are addressing these challenges.

As Editor-in-Chief, I am particularly gratified by the attention that *IoT* is acquiring, as evidenced by the number of proposals for special issues that have been submitted. We are adapting the Editorial Calendar to accommodate the proposals. The *IoT* Editorial Calendar for the next 24 months is listed in the accompanying table. Individuals and organizations involved in the topic areas listed are invited to submit articles.

The Calls for Special Issues can be found on the *IoT* website: <https://www.comsoc.org/publications/magazines/iee-internet-things-magazine>. Articles not specifically addressing the topic areas will also be accepted. The *IoT* General Call for Articles can be found at <https://www.comsoc.org/publications/magazines/iee-internet-things-magazine/cfp/general-call-articles> and is reproduced below.

CALL FOR ARTICLES

The *Internet of Things Magazine (IoT)* publishes high-quality articles on IoT technology and end-to-end IoT solutions. *IoT* articles are written by and for practitioners and researchers interested in practice and applications, and selected to represent the depth and breadth of the state of the art. The technical focus of *IoT* is the multi-disciplinary, systems nature of IoT solutions. *IoT* is a forum for practitioners to share experiences, develop best practices, and establish guiding principles for technical, operational and business success.

The magazine is currently soliciting articles for publication. Articles should examine one or more actual deployments of an IoT solution and discuss:

- A high-level operational description of the IoT solution, addressing: the problem space; a summary of systems operation; and how the overall problems were solved.
- A high-level technical description of the IoT system: What technical challenges were encountered? What solutions were developed? What were the technical risks encountered in development? How were they overcome?
- A summary of the business case: What kind of benefits did the stakeholders receive from the solution? Were they greater than or less than expected? Were any policy or regulatory issues encountered?
- Lessons learned from deployment and operation: What were the key lessons learned? Can this experience contribute to defining best practices? What were the risks and rewards?

Articles should be general and present real-world experiences, with the intended audience being all members of the IoT com-

Topic	Submission due date	Publication date
Communications for IoT: Connectivity and Networking		March 2020
Blockchain-Enabled Industrial IoT		June 2020
Computational Intelligence in Internet of Everything	February 2020	September 2020
IoT and the Environment	April 2020	December 2020
IoT and Privacy	June 2020	March 2021
IoT and eHealth	August 2020	June 2021
IoT and Defense	November 2020	September 2021
Standards for IoT	March 2021	December 2021

munity, independent of technical or business specialty. Articles are expected to add to the knowledge base or best practices of the IoT community; sales/marketing materials are not appropriate. Authors are asked to strive to make their papers understandable by the general IoT practitioner. Mathematical material should be avoided; instead, references to papers containing the relevant mathematics should be provided. Authors are encouraged to use color figures and submit multimedia material along with their articles for review. Authors should target 4,500 words or less (from introduction through conclusions, excluding figures, tables, and captions), or six (6) pages. Figures and tables should be limited to a combined total of six. The number of archival references is recommended not to exceed fifteen (15).

IoT also publishes regular columns on topics of interest to IoT practitioners. Topical columns update readers on issues and events in the world of IoT. Regular columns are published in the following areas:

- Around the World of IoT — Recent events or technology developments in IoT.
- Bridging the Physical, the Digital, and the Social — Socially-aware advancements in IoT.
- Policy and Regulatory Affairs — Discussions and reports on policy issues facing the world of IoT.
- IoT Standards — Discussions and reports on efforts in standardization of IoT technology and systems.
- Privacy and Security — Discussions and reports on interaction of IoT with privacy and security concerns.
- Book Reviews

Columns should be of general interest to all members of the IoT community. Columns should inform the reader about issues and events that may affect the business and practice of IoT; sales/marketing materials are not appropriate. Authors are asked to strive to make their articles understandable by the general IoT practitioner. Authors should target 1500 words or less (from introduction through conclusions, excluding figures, tables, and captions), or two (2) pages. Figures and tables should be limited to a combined total of two. The number of archival references is recommended not to exceed five (5).

Authors should submit articles and columns to <https://mc.manuscriptcentral.com/iotmag>.

IoT does not have a specific template and does not require manuscripts to be submitted in any specific layout. However, authors can use the template for *IEEE Transactions* to get a rough estimate of the page count: https://www.ieee.org/publications_standards/publications/authors/author_templates.html.

Advance Your Career with IEEE ComSoc Communications Technology Courses

IEEE ComSoc offers high-quality online training courses taught by industry experts and designed to deliver you the skills you need to advance your communications technology career. Plus, you can earn IEEE Continuing Education Units (CEUs) with each course.

UPCOMING TRAINING COURSES

4 March

5G New Radio Non-Standalone Deployment: Implementation, Performance and Challenges

Learn the fundamentals of the architecture, air interface and 3GPP deployment options of 5G New Radio (5G NR) as well as the interworking of NR with LTE regarding radio functionality, bearer handling, mobility and performance.

11 March

Wireless Positioning and Location

Study positioning fundamentals and location methods as well as current issues including: GNSS, TOA and TDOA, RSS, indoor location, inertial positioning and data fusion.

18 March

Wireless for the Internet of Things

Build the skills to create products for use with IoT applications, regardless of chosen platform.

25 March

Fundamentals of 5G Small Cell Deployments

Explore 5G small cells architecture options, their effect on traffic engineering and how they will change the way the operator supports the network.

TRAIN YOUR TEAM!

IEEE ComSoc Training offers discounts for group course registrations of eight or more seats. Plus, we can work with you to develop a custom, online or in-person training experience to match the unique needs of your group.

Find out how we can help your team meet its professional development and continuing education needs. Contact Tara McNally at t.mcnally@comsoc.org for more information.

DISCOUNTS AVAILABLE FOR IEEE COMSOC MEMBERS!

Visit www.comsoc.org/training today to register or learn more.

Join the world's largest technical
professional organization!
Save 50% Off on Membership!



Networking • Conference Discounts • Technical Publications • Volunteer



IEEE ComSoc Affiliate Membership* Only \$54.00!

Valuable Discounts on IEEE ComSoc Conferences

ComSoc members save an average of \$200 on ComSoc-sponsored conferences.

Free Subscriptions to Highly Ranked Publications**

Receive digital access to IEEE Communications Magazine, IEEE Communications Surveys and Tutorials, IEEE Journal of Lightwave Technology, IEEE/OSA Journal of Optical Communications and Networking and many other publications – every month!

*Society affiliates are members of a Society only and are not entitled to any IEEE benefits or services that are reserved solely for IEEE members.

**2017 Journal Citation Reports (JCR)

IEEE WCET Certification Program

Grow your career and gain valuable knowledge by completing this certification program. ComSoc members can save \$100.

IEEE ComSoc Training Courses

Learn from industry experts and earn IEEE Continuing Education Units (CEUs) / Professional Development Hours (PDHs). ComSoc members can save over \$80.

Exclusive Events in Emerging Technologies

Attend events held around the world on 5G, IoT, Fog Computing, SDN and more! ComSoc members can save over \$60.

Visit our website to join today and learn more!

<https://www.comsoc.org/>

IEEE Foundation

REALIZE THE FULL POTENTIAL OF IEEE



ILLUMINATE

the possibilities of technology by using it to address global challenges



EDUCATE

the next generation of innovators and engineers



ENGAGE

a wider audience in appreciating the value and importance of engineering and technology



ENERGIZE

innovation by celebrating technological excellence

The world's most daunting challenges require innovations in engineering, and IEEE is committed to finding the solutions.

The IEEE Foundation is leading a special campaign to raise awareness, create partnerships, and generate financial resources needed to combat these global challenges.

Our goal is to raise \$30 million by 2020.

DONATE NOW
ieeefoundation.org





Introducing IEEE Collabratec™

The premier networking and collaboration site for technology professionals around the world.

IEEE Collabratec is a new, integrated online community where IEEE members, researchers, authors, and technology professionals with similar fields of interest can **network** and **collaborate**, as well as **create** and manage content.

Featuring a suite of powerful online networking and collaboration tools, IEEE Collabratec allows you to connect according to geographic location, technical interests, or career pursuits.

You can also create and share a professional identity that showcases key accomplishments and participate in groups focused around mutual interests, actively learning from and contributing to knowledgeable communities. All in one place!

Network.
Collaborate.
Create.

Learn about IEEE Collabratec at
ieeecollabratec.org