Cloud computing to automate phenotype collection and data analyses in dairy systems

R. E. P. Ferreira[1] and J. R. R. Dórea[1,2,*]
[1]*Department of Animal and Dairy Sciences, University of Wisconsin-Madison, Madison, WI 53706, USA*
[2]*Department of Biological Systems Engineering, University of Wisconsin-Madison, Madison, WI 53706, USA*
[*]Corresponding author: João R. R. Dórea, joao.dorea@wisc.edu

Abstract

Among technologies used for animal high-throughput phenotyping, computer vision generates large amounts of unstructured data. The development of automated data pipelines using scalable solutions such as cloud computing, can be an effective strategy to generate animal-level information in real-time. The main objective of this study was to develop a cloud computing-based system to automate phenotype extraction in dairy farms using computer vision systems. In this study we used animal identification and body condition score (BCS) as case studies to demonstrate the automation pipelines for data collection and analyses. The proposed system outputs a predicted BCS for each cow individually each day, using depth and infrared images captured from cameras installed at the exit of the milking parlor and uploaded to the cloud. The entire pipeline is performed on the cloud and consists of four different steps: (1) determine good quality images; (2) remove pixels that do not belong to the cow's body; (3) identify the cow based on coat color patterns; (4) and assess BCS using depth images. The neural networks designed to identify individual cows and score BCS achieved accuracies of 94% and 71%, respectively, on a group of 59 cows, with a mean absolute error of 0.20 for BCS. From a total of 52,247 snapshots, 32,655 were discarded in step 1, and the remaining 19,592 were evaluated. All predicted BCS values and the associated animal identification (ID) are automatically stored in a relational database once a new snapshot is captured and sent to the cloud. The main contribution of this work is the development of a scalable cloud-based framework for large-scale individual phenotyping in livestock systems using computer vision techniques.

Keywords: cloud computing, computer vision, dairy cattle, body condition score

Introduction

As global food demand increases with the global population, food production systems are pressured to become more cost-effective and efficient, both production- and environment-wise, in order to satisfy such increasing demand. In livestock systems, such efficiency can be achieved mainly through better management practices and genetic improvement. Both of these routes can greatly benefit from high-throughput phenotyping, which consists of measuring and monitoring key traits from living organisms, ideally in a non-invasive, automated, and scalable manner (Koltes et al., 2019). High-throughput phenotyping, more specifically through precision livestock farming (PLF) technologies, allows for better informed and rapid farm management decisions, as well as accelerated genetic selection capabilities (Berckmans, 2017; Brito et al., 2020).

Precision livestock technologies provide a great way to achieve high-throughput phenotyping, with computer vision-based systems emerging as potential approaches that provide non-invasive, automated, and scalable solutions for individual animal monitoring (Fernandes et al., 2020). However, a great challenge in applying PLF technologies in livestock farms pertains to making efficient use of the generated data (Koltes et al., 2019). As data generated by each PLF system are usually available only locally in the farm or even locked within each provider's software, data integration and availability become exceedingly difficult, and the potential for new solutions that could benefit from multiple data sources remains unexplored.

Cloud computing technology could be used to store the data generated by PLF systems available on the farm and process such data into valuable information for the farmer, which could be accessed from anywhere with internet connection. Integrating PLF technologies into cloud computing solutions can mitigate the problems related to data integration and availability, as all data generated from different sources are stored and made available on a single platform. Furthermore, cloud computing allows for great flexibility in scaling up data analyses, storage, and other types of data processing. For example, machine learning algorithms that process the collected data into useful information can be developed and seamlessly deployed to the cloud infrastructure to process few or billions of data points as it becomes available by the data collection systems. In this study, we propose a fully automated framework for processing data collected through image sensing devices located on farms into individual animal phenotype and information through a cloud computing infrastructure. We demonstrate the proposed framework using animal identification and body condition scoring (BCS) as case studies, as individual animal identification is frequently considered a first step required for further animal phenotyping, and BCS is an important metric for assessing a dairy herd and guiding management decisions (Roche et al., 2009).

Materials and methods

Data collection

Images from 59 lactating cows housed at the Emmons Blaine Dairy Cattle Research Center (Arlington, WI) were taken using four Intel RealSense D435 depth cameras installed at the milking parlor exit lanes. The cameras were arranged to capture top-view images of the cows exiting the milking parlor twice a day, during milking scheduled times, at a rate of 4 snapshots per second. Each snapshot contained a depth and an infrared image, both with a resolution of 640 x 480 pixels. The depth images contain, for each pixel, a value that represents the distance from the object in that pixel to the camera, in millimeters. The infrared images contain, for each pixel, a value between 0 and 255, ranging from plain black to plain white, respectively, and they were used in this study to identify cows through their coat color pattern using a computer vision algorithm deployed in the cloud platform. Examples of a depth and an infrared image are shown in Figure 1.

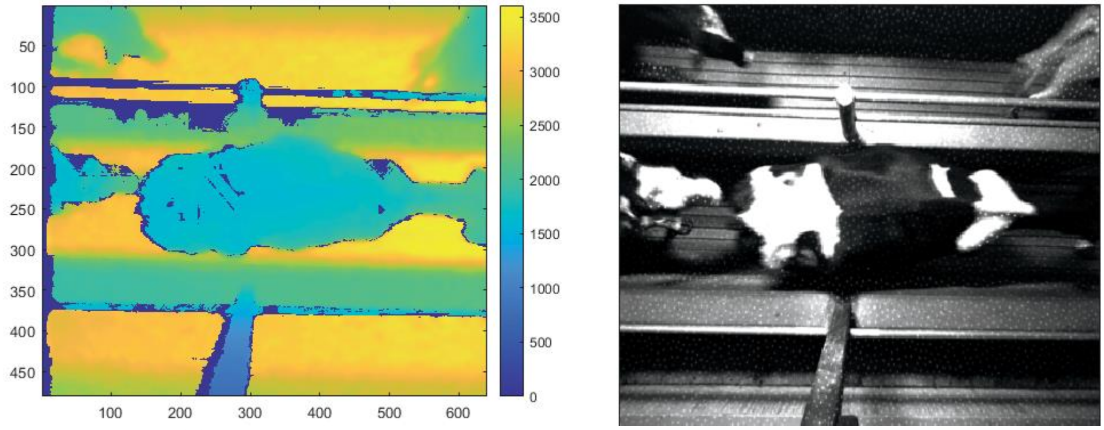

(a) Depth image (b) Infrared image

Figure 1: Example of a snapshot containing a depth (a) and an infrared image (b). In a depth image, each pixel contains a value that represents the distance in millimeters from the object in that pixel to the camera sensor. In an infrared image, each pixel contains a value that represents the light intensity captured by the infrared sensor in that pixel's location.

Data analysis framework

In order to predict the BCS of an individual cow from a snapshot, four steps were performed for each depth and infrared pair: an image classifier (Classifier1) inferred, using the depth image, whether that snapshot contained a poorly positioned cow, and in those cases the snapshot was removed from the dataset; an image segmentation algorithm (Segmenter) used the depth image to detect all pixels containing the cow's body, and the resulting mask was used to segment both the depth and infrared images; an image classifier (Classifier2) identified which cow was in the snapshot using the segmented infrared image; and an image classifier (Classifier3) scored the cow's body condition using the segmented depth image. Figure 2 shows examples of a segmentation mask resulting from the Segmenter algorithm and the corresponding segmented infrared and depth images.

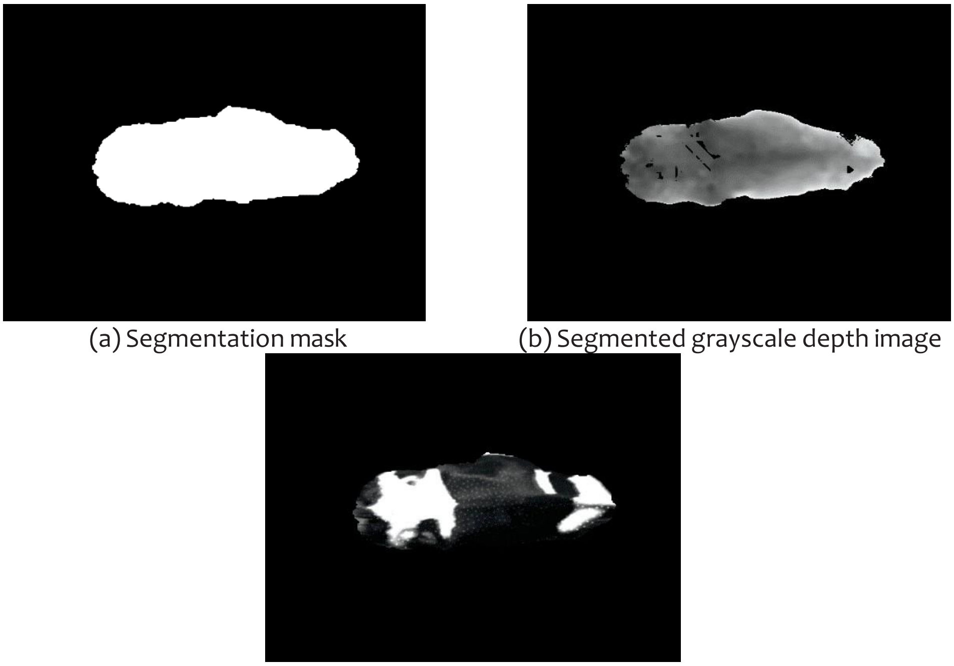

(a) Segmentation mask (b) Segmented grayscale depth image

(c) Segmented infrared image

Figure 2: Examples of a segmentation mask (a) and the corresponding segmented depth (b) and infrared (c) images. The segmentation mask results from the Segmenter algorithm after receiving a depth image (1a) as input, containing 1 (white) for pixels in the cow's body, and 0 (black) otherwise. The mask is then applied to the depth (1a) and infrared (1b) images, resulting in segmented depth (b) and infrared (c) images, respectively.

All algorithms were trained separately and deployed to a cloud-computing environment hosted in Microsoft Azure, following a modular microservices architecture. Microsoft Azure hosts the developed code that runs predictions on new images using the trained algorithms, callable through a Representational state transfer (REST) Application programming interface (API). The three image classifier algorithms were implemented as deep neural networks based on the Xception architecture (Chollet, 2017), and the image segmentation algorithm was implemented as a deep neural network based on the U-Net architecture (Ronneberger et al., 2015). These four algorithms were deployed as their own separate services, each receiving an image as input and the desired information as output. The Classifier1 service receives a depth image and outputs the probability that the snapshot contains a well-positioned cow. A cow is considered well-positioned if its whole body from tail to neck can be seen in the image without any occlusion. The Segmenter service receives a depth image and outputs an image of the same resolution as the received image, with each pixel holding a value of 1 if it is contained in the cow's body (from tail to neck, excluding the head), or 0 if it is part of the background. The Classifier2 service receives a segmented infrared image and outputs two values: the identification number (ID) of the cow with the highest confidence of being the one in the image, and the corresponding confidence value. Finally, the Classifier3 service receives a segmented depth image and outputs an array of probabilities that the cow contained in that image has each of the possible body condition scores, ranging from 1 to 5, in quarter point increments. An orchestrator service was then responsible for calling each of the four services in the correct order and saving the results in a Microsoft SQL Server database for further analysis. This orchestrator service is triggered every time a new image is uploaded to the corresponding Microsoft Azure Blob Storage. All five services (four deep neural networks and the orchestrator) can be called independently via REST API, further increasing the framework's flexibility and potential for integration with other systems. Using this modular approach allows for re-use of certain core services, such as body segmentation, and facilitates the implementation and deployment of new functionalities into the cloud platform, all seamlessly to the farm operations as no updates are required in the on-premises farm computer infrastructure.

Since cows were milked twice a day in the farm used in this project, the developed system could assign a daily BCS for each cow. However, for the purpose of evaluating the scoring quality, we compared the BCS outputted by the system with the score assigned by trained humans in four different time points, as it would be too time consuming to have humans manually assign BCS for each of the 59 cows every day. The time points were defined to roughly follow two-week intervals: August 12th, August 25th, September 15th, and October 6th, 2020. Each snapshot was assigned a BCS and a cow ID. The BCS was calculated as the weighted average of scores, with the weights being the probabilities output from Classifier3, and the cow ID was the output from Classifier2, with a corresponding confidence value. To obtain a final BCS for a given cow each day, all snapshots were grouped based on the date and cow ID, and a simple average BCS was calculated. To avoid using incorrectly classified images, only images with a cow ID confidence greater than 95% were used. This final BCS for each cow and day was then compared to the average of values assigned by three trained humans to evaluate the system accuracy. For that, the mean absolute error (MAE) and root mean squared error (RMSE) metrics were calculated. Images from August 8th to August 12th were used to train the algorithms, and the images taken on the other three time points (August 25th, September 15th, and October 6th) were used to evaluate the system's performance. For the BCS classifier (Classifier3), the quarter point that was the closest to the average of the three manually assigned scores was used as label.

Results and discussion

From a total of 52,247 snapshots, 32,655 were discarded as they were classified as not adequate by Classifier1, and the remaining 19,592 were further evaluated. This represents a reduction of 62.5% in the total image data stored in the system, keeping just the data that is relevant for the current case study. It is important to highlight that Classifier1 could also be implemented in edge computing systems and perform similar image filtering at the edge device before further processing. The service performing cow identification (Classifier2) achieved an accuracy of 94% for classifying the cow present in each image. This accuracy was calculated using a holdout set of the training data containing only images taken on August 12th, after training using images taken on August 8th.

To evaluate our proposed system and compare it with the accuracy of BCS assigned by trained humans, we defined the ground truth BCS as the average of the three human-assigned scores for each cow and each day. Table 1 shows the calculated root RMSE, MAE, and coefficient of determination (R^2 score) for each human BCS assigner and the proposed system. It is worth noting that performance from human assigners might be overestimated as the ground truth values were defined as an average of their assessments.

Table 1: Comparison of each BCS assigner performance using RMSE, MAE, and R^2 score as evaluation metrics. The proposed system achieved comparable performance to human assigner 3.

BCS assigner	RMSE	MAE	R^2 score
Human 1	0.170	0.130	0.86
Human 2	0.178	0.141	0.86
Human 3	0.293	0.229	0.78
Proposed system	0.262	0.204	0.54

It is noticeable that our system performs well in respect to RMSE and MAE, ranking between the human trainers and above the naïve baseline. However, it performs worse than all human trainers when considering R^2 score. Looking in more detail at the predicted values, it is noticeable that the proposed system performs poorly in cows with a high BCS value, as shown in Figure 3. This is possibly due to imbalances in the training dataset, with more examples of cows with intermediate BCS values (between 2.75 and 3.75) than with higher BCS (above 4.0). Because of that, the trained BCS classifier network tended to predict intermediate values even for actual high BCS cows, and this is an issue commonly experienced when working with imbalanced datasets, and particularly in automatic BCS evaluation systems (Qiao et al., 2021). There exist different techniques that could be used to address this issue, such as training using a weighted loss (less common classes get higher weights so the network can learn about them more quickly). Other techniques consist of downsampling or upsampling the dataset, meaning we could either remove examples of more common classes, or artificially add examples of less common classes by using a technique called image augmentation (Perez et al., 2017).

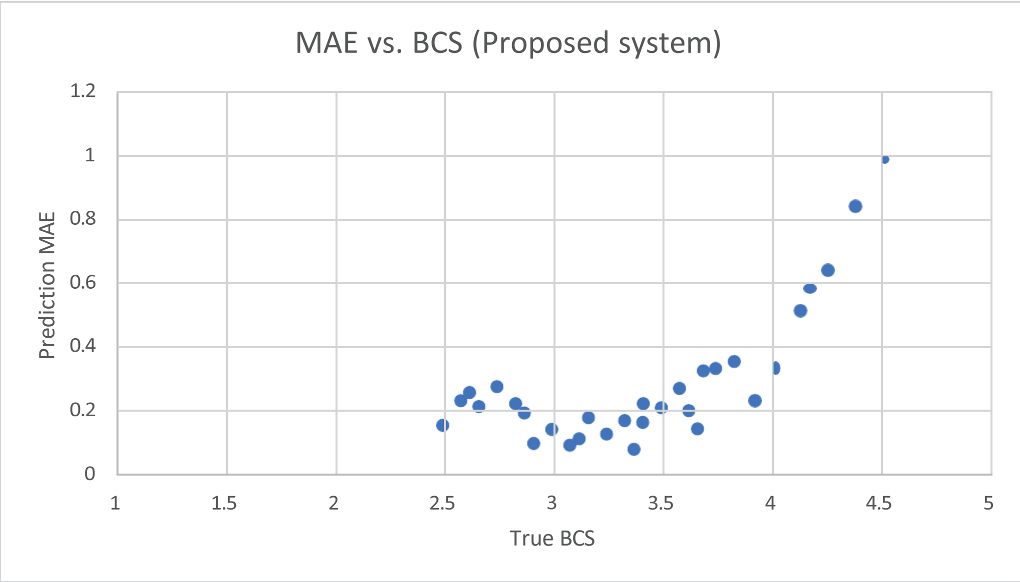

Figure 3: MAE vs. BCS plot for the proposed system. Each point is a pair of true BCS and the corresponding MAE of the predictions. Equal values of BCS were grouped together and a single average MAE was calculated.

Finally, it is worth noting that the main objective of this study was not to develop the most accurate and robust BCS classifier, but to demonstrate the potential of the proposed cloud computing-based framework in a case study that represents a real-world problem in dairy farms. This framework allows for automatic processing, inference, and summarization of the data collected by systems available on the farm (in this case, depth sensing cameras). Since all collected data as well as the algorithms' results are stored in the cloud, this data can be accessible in real-time through reports available in the cloud, generating valuable insights to interested parties without the need to be physically present at the farm. An example of such an automatically generated report is shown in Figure 4.

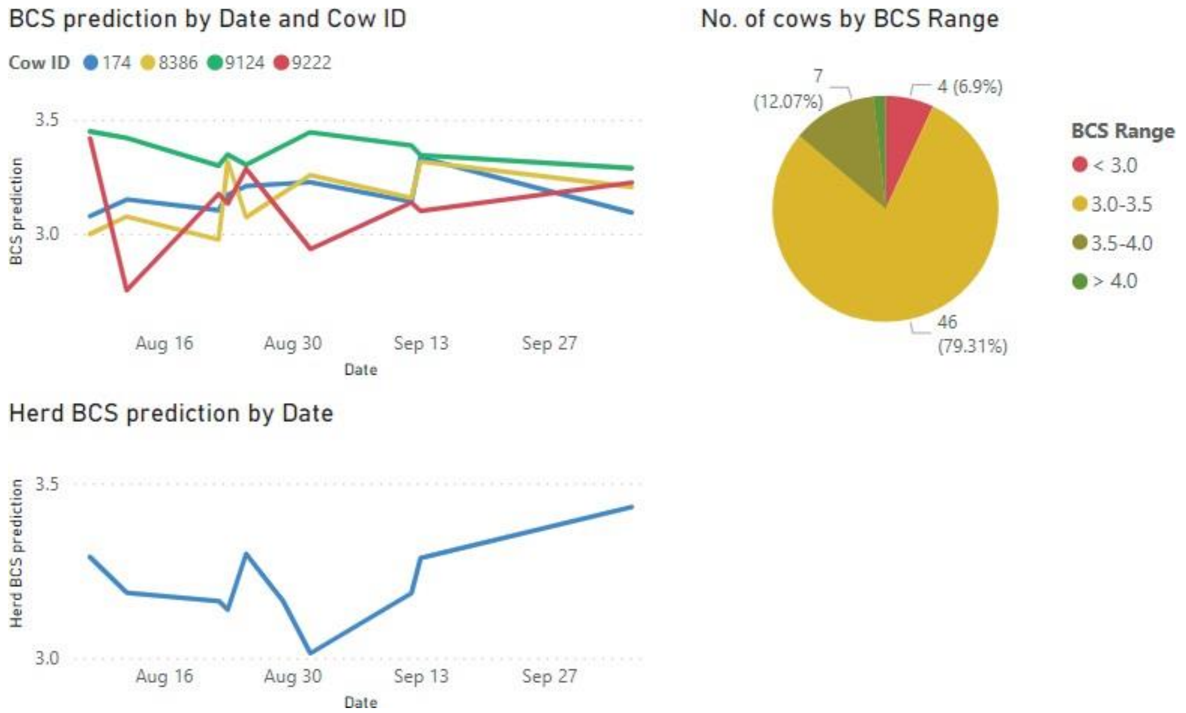

Figure 4: Example of an automatically generated report of BCS for a dairy farm, using data available in the cloud resulting from cow ID and BCS computer vision algorithms. The top-left chart shows the individual evolution of BCS for 4 different cows, the bottom-left show the evolution of the average herd BCS, and the top-right chart shows the current proportion of cows in each BCS range.

Although internet connection is becoming widely available in farms across the world, the connection is still often not reliable enough to transfer large amounts of data with good speed. Because of that, transferring all data to the cloud might not be optimal or even feasible in certain scenarios. Under those circumstances, edge computing can become a potential solution that relies on processing most of the collected data locally and sending just essential information to the cloud (Shi et al., 2016). However, such systems have the disadvantages of requiring powerful hardware available at the farm, and permanently losing the raw collected data, which might not be desirable in certain applications. It is imperative that both alternatives are evaluated when deciding on implementing a PLF system in a farm, evaluating whether the available internet connection is reliable enough to transfer the required information to the cloud, whether raw data storage is required, and the local hardware capabilities.

Conclusions

This study proposed a cloud computing-based framework for processing data collected through image sensing devices located on farms into individual animal phenotype using computer vision techniques. In the demonstrated case study of cow ID and BCS, our proposed system performed similarly to human evaluators, meaning that it has the potential to be used in real farm operations as an alternative to automate and facilitate BCS evaluation processes, and motivate more farms to incorporate BCS monitoring into their management routines. It is important to note that the final BCS estimation accumulates errors from all previous steps, meaning that errors in image quality assessment, cow identification, and body segmentation can be carried over to the final BCS prediction. Therefore, each neural network implemented in the pipeline should be carefully evaluated independently, as tuning and re-training can be done in case the system performs poorly in different contexts such as a new farm or research facility. The modular nature of the proposed approach allows for such reimplementation and flexibility, enabling certain core services to be re- used, such as body segmentation, and facilitating the implementation and deployment of new or enhanced functionalities into the cloud platform. New applications and upgrades can be deployed seamlessly to the farm operations as no updates are required in the on-premises farm computer infrastructure. Evaluating and enhancing separate parts of the system after the proposed framework is implemented is key for the successful deployment of PLF systems in commercial farms.

Acknowledgments

This project was funded by the USDA National Institute of Food and Agriculture (Washington, DC; grant 2020- 67015-30831) and Hatch project (WIS03085).

References

Berckmans, D. (2017) General introduction to precision livestock farming. *Animal Frontiers* 7(1), 6-11.

Brito, L.F., Oliveira, H.R., McConn, B.R., Schinckel, A.P., Arrazola, A., Marchant-Forde, J.N., and Johnson, J.S. (2020) Large-scale phenotyping of livestock welfare in commercial production systems: A new frontier in animal breeding. *Frontiers in Genetics* 11, 793.

Chollet, F. (2017) Xception: Deep learning with depthwise separable convolutions. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition,* 1251-1258.

Fernandes, A.F.A., Dórea, J.R.R., and Rosa, G.J.D.M. (2020) Image analysis and computer vision applications in animal sciences: an overview. *Frontiers in Veterinary Science* 7, 551269.

Koltes, J.E., Cole, J.B., Clemmens, R., Dilger, R.N., Kramer, L.M., Lunney, J.K., McCue, M.E., McKay, S.D., Mateescu, R.G., and Murdoch, B.M. (2019) A vision for development and utilization of high-throughput phenotyping and big data analytics in livestock. *Frontiers in Genetics* 10, 1197.

Perez, L., and Wang, J. (2017) The effectiveness of data augmentation in image classification using deep learning. In: *arXiv preprint arXiv:1712.04621.*

Qiao, Y., Kong, H., Clark, C., Lomax, S., Su, D., Eiffert, S., and Sukkarieh, S. (2021) Intelligent perception for cattle monitoring: A review for cattle identification, body condition score evaluation, and weight estimation. *Computers and Electronics in Agriculture* 185, 106143.

Roche, J.R., Friggens, N.C., Kay, J.K., Fisher, M.W., Stafford, K.J., and Berry, D.P. (2009) Invited review: Body condition score and its association with dairy cow productivity, health, and welfare. *Journal of dairy science* 92(12), 5769-5801.

Ronneberger, O., Fischer, P., and Brox, T. (2015) U-net: Convolutional networks for biomedical image segmentation. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention* Munich, Germany, 234-241.

Shi, W., Cao, J., Zhang, Q., Li, Y., and Xu, L. (2016) Edge computing: Vision and challenges. *IEEE Internet of Things Journal* 3(5), 637-646.