CIDAI Centre of Innovation for Data tech and Artificial Intelligence

Smart Farming: Estimating weight from PIGtures

Coen Antens
Marc Maldonado
Juan Ramón Jiménez
Richard Segovia Barreales

CVC Computer Vision Center

Suggested citation. CIDAI-POC-2022-07 // Using Artificial Intelligence to estimate the weight of a pig by Coen Antens, Marc Maldonado, Juan Ramón Jiménez, Richard Segovia Barreales – CVC, CIDAI, 2022. Licence: CC BY-NC-SA 4.0

1. Abstract

1.1. Objectives

Over the years the agricultural sector has evolved along with new emerging technologies, a concept that is known known as Smart Farming. This new idea takes advantage of the potential of ICT to reduce the cost of both production costs and farm efficiency. Based on this idea, new management models have been devised, ranging from intelligent greenhouses and managing diseases, towards the implementation of agricultural drones that by using aerial images allow significant savings of time and labor at the time of visual verification of a crop. In addition, this new idea is not only applicable to cultivation, but that there are also interesting applications in cattle.

In this context, CVC worked on a Smart Farming project related to the automation of a pig fame, where one of the main challenges is to estimate the weight of the pigs with the motivation of knowing their evolution over time and prepare them to obtain the optimum weight when they arrive to the slaughterhouse and thus obtain the maximum possible return. At the time of delivering a pig to the slaughterhouse its value depends on how how close the weight is to the ideal weight of 120kg. The price that will be paid for the animal decreases when the weight is below or above the standard established and that is why the control of the evolution of this animals is of significant economic interest.

Índex

1.2. Summary of the problem

When fattening pigs it is very important to verify the evolution of theirs weights over time. Fattening pigs arrive at the pig farm when they weigh about 20 kg. Next, about 4 months later they weigh about 120 kg and that's when they have to be transported to the abattoir. As the prize that the slaughterhouse will pay for the pig depends on how close the actual weight approximates the value of 120 kg, the farmers has a special interest in monitoring the weight of the pigs over time.

However, weighing a pig is very, very time consuming because it has to be done by hand with a scale that is like a cage or box. The pigs have to be introduced in this cage in order to weigh them. Of course, introducing the pigs one by one into the cage is not an option, because the pigs are not willing to corporate. Therefore, instead of weighing the pigs directly, the idea is to weigh them indirectly by means of a 3D camera. With a 3D camera an approximation of the volume of the animal can be obtained and the volume is related to the weight of the pig. Having a lot of examples of depth maps and the corresponding weight, a system can be trained to estimate the weight of a pig using a depth image.

1.3. Solution

Analyzing the context, the distribution of our farm and taking into account the daily routine of the animals, a 3D camera was integrated in a traditional scale that is used for weighing the pigs. This camera is an industrial camera, IP67 and thus resistant to special circumstances of a pig farm (Figure 1). The actual weight of the pig is still obtained from the scale. However, for the project the scale has incorporated a number of photosensitive cells that make sure that only one pig enters the scale. In addition, a RFID reader reads the chip of the pig as soon as the photosensitive cells detect the pig that enters the scale. After weighing the animals, the weight and the identifier of the pigs are sent to a PLC and finally a trigger is generated for the 3D camera to take a picture. In this way, a mapping is established between the weight of the pig and the corresponding image (depth map).

Figure 1. 3D Camera to monitor the pigs

With this system setup a lot of data is generated automatically. Next, the idea of the project is to develop a mathematical model based on Deep to eliminate the costly process of using a scale. The current setup will be used to create a dataset and to train the model. After training the model, the information captured by the 3D camera will be enough to come up with an estimation of the weight of the pig. Thus, the weight of the individual pigs can be monitored over time. On the one hand this information can be used to verify whether the weight of the pig is evolving well and on the other hand to automatically detect when the weight of the pig is optimal.

2. Description of the problem you are solving

The pig sector is a very important sector in Catalonia. According to the Statistical Institute of Catalonia, the pig sector with about 6,500,000 animals occupies the second place in the ranking after the chickens with more than 43,000,000 animals. A large part of the pigs are fattening pigs and for this reason a framework was chosen to experiment with Smart Farming techniques. The growth of pigs is carried out in fattening farms. Pig fattening lasts about 4 months. The final weight of the pigs is 110-120 kg and they are taken to the slaughterhouse to scarify them. In this project, various aspects of pig fattening have been studied and a set of sensors (cameras) have been installed both on the farm and in the slaughterhouse to obtain data that help evaluate the process, monitor it and improve it (Figure 2).

A very important factor in the pig fattening process is its feeding. The classic way to feed pigs is with feed. But, there are some trends that are experimenting with other types of food, such as liquid food that is easier for animals to digest. There are even initiatives that contemplate the introduction of this type of food to farms. So, it has been considered interesting to define a project with different actors who have a relevant role in the process around the life of the pigs that are raised for meat consumption.

We have seen before that when delivering a pig to the slaughterhouse its value depends on how much its weight approximates the ideal weight of 120kg. The price that will be paid for the animal decreases when the weight is below or above the standard established that is why the control of the evolution of the pigs is so important.

Figure 2. Overview of the project

3. Implementació de la solució

3.1. Architecture, technology and data used

3.1.1 Architecture

The first option was to implement a background subtraction idea, because at the beginning of the project there were not enough images to be able to crop the pig using Machine Learning. As there was no image of the empty scale available at the beginning to be able to perform a background subtraction, an attempt was made by averaging all the images to generate a new one which would be used as background (similar techniques are applied to obtain a photo of for example a Egyptian Pyramid without any tourists). We experimented with this approach both with both infrared camera images and depth images. Once this background image was generated, a subtraction is performed on each image to eliminate the background and then we applied some Gaussian smoothing to achieve a cleaner image. When generating the final masks we found that the results obtained with the infrared images are generally noticeably better than the results obtained with the depth images.

After a while, a background image was provided to perform the subtraction. This time, instead of taking the average of all the images, we perform the subtraction directly with the background. In the same way, a smoothing filter is applied and applying some thresholds the new masks are generated. These techniques obtain a performance similar to the two previous techniques, only the other way around. Regarding the use of infrared images, this result is due to the fact that now, unlike the previous one, we use an image of the bottom and due to the surface of the scale a reflection of light similar to that of the pig's back is generated, so most pigs are segmented throughout their body except for the part that coincides with the reflection and this explains the loss of precision.

The improvement using depth images with respect to the previous method using the same is due to the fact that there is a difference in the background of the image. This is where we realized that one of the walls of the scale changed position depending on the photograph. This is because the walls can be adjusted to the size of the animals to make sure that no more than one pig enters the scale. Even in this way, of all the available images there is a large number with the wall in the same position with respect to other images with the wall in a different position. As the background image is a representative of the images that are repeated the most in terms of the position of the wall that is why it generates such good results.

At this point it clear that the real problem has to do with the bottom of the scale and not with the pigs so we decided to look for another solution. The first idea was to use Canny and Hough to detect the walls automatically and thus eliminate the problem of moving walls and bottom. This technique worked well with the first set of images, but after receiving new images it became an impossible challenge due to the variety of photos.

Object Detection

Taking into account the potential of cropping the image, we decided to generate a Ground Truth to train an object detection mo-

del. So, this meant that we had to manually segment a large number of images. For this the Hasty.ai tool was used, which provides various tools for labeling images both for semantic segmentation and object detection. In this case, we chose to use semantic segmentation, since from here we can also generate bounding boxes for object detection.

Figure 3. Objection detection

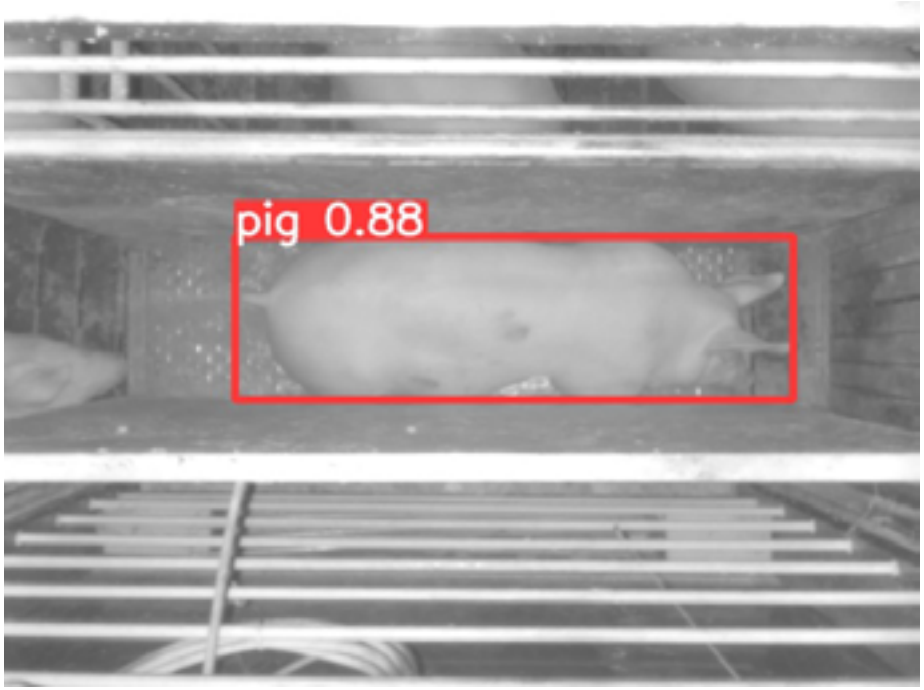

The tool allows labeling simultaneously with other people and during tagging it trains models that offer a possible mask which you can edit. With this Ground Truth the bounding boxes were extracted and a Yolo model was trained, in particular YOLO V5.

We use Yolo's detection to find the pig in the image and the pig is cut out using a threshold that was picked by Otsu. It generates a result similar to results explained in the previous session (using classical Machine Vision), but slightly worse in terms of accuracy (Figure 3).

We chose the Yolo approach anyway because the difference is negligible and it ensures that we always catch the important pig in the image and no other pigs that poke their heads through the door as can be seen in some image. In addition, morphology techniques are applied to these generated masks to fill gaps inside the pigs generated by their stains and horizontal kernels to try to horizontally dilate the mask and recover non-segmented points belonging to the pig.

Semantic Segmentation

A part from experimenting with Object Detection networks like the YOLO network we decided to consider the problem as a Semantic Segmentation problem. In order to find a better solution for the segmentation problem, different architectures were set up. First we trained a standard U-Net model with ResNet-50 backbone. But we also considered lighter models such as MobileNetV2.

Finally, after different tests with different backbones and different hyperparameter configurations, a model was obtained that, trained with the images captured by the infrared camera and the masks generated with HastyAI, obtained an Intersection Over Union of 0.98 using an InceptionV3 backbone. In the process we used the images that we segmented manually to make sure that

only real information is taken into around. Only this more relevant information will be then be used in the regression model and will also speed up the training of the model. Later in the development we decided to use only the depth images to train the segmentation model. Not only that, but we also decided to refine the masks by eliminating the pig's head. This was done because sometimes the pig's head is moving around a lot and thus may introduce some extra difficulties when estimation the pig's weight.

Other transformations

We have described the way in which we segmented the pigs in order to reduce the information and facilitate the regression task. Under the premise of regularizing the data, a transformation is performed to center the pig in the center of the image, that is, the midpoint of the cropped mask is calculated and the image is moved to center it.

Finally, with the idea of using GNNs for weight estimation, a 3D model of the pig is built by applying the mask generated with the semantic segmentation and considering only the points in 3D that corresponds with this mask. With Open3D we refined the model eliminating points that corresponded to the bottom of the scale and other outliers and unwanted points were eliminated using a statistical approach.

Figure 4. Example of a point cloud corresponding to a pig

These point clouds were treated to transform them into a graph in the form of triangular meshes. For the task of building the meshes, the Open3D library was used. This library provides different functions for surface reconstruction: Alpha shapes, ball pivoting} and Poisson surface reconstruction. By adjusting the parameters of the Poisson function finally we obtained the reconstruction shown in Figure 4. With these point clouds and these generated meshes, an attempt was made to find a solution to the regression problem, but it was ruled out due to its complexity and the scarce development of these technologies in these areas.

3.1.2 Data used

Before we started to build the regression model, we first analyzed the data. The data provided using the setup of Figure 5 was the following:

- Top view infrared images of the pig in the scale
- Top view depth images of the pig in the scale
- XLS with weight information, the ID of the pig, time and date of weighing

The first task was to proceed to map the images with the corresponding weights in the XLSX file. Each image consists of a series of digits in the name that indicates the timestamp of when the picture was taken. These instants of time did not fit with the time that marked the entry of the weights in the Excel file: there was a difference of 4 min delay. After correcting this information, a CSV file was generated taking into account this time difference. After this process, some images were discarded of which there was no entry in the Excel file and all images before a certain date were deleted because we decide to change the way in we capture the images.

Figure 5. System setup

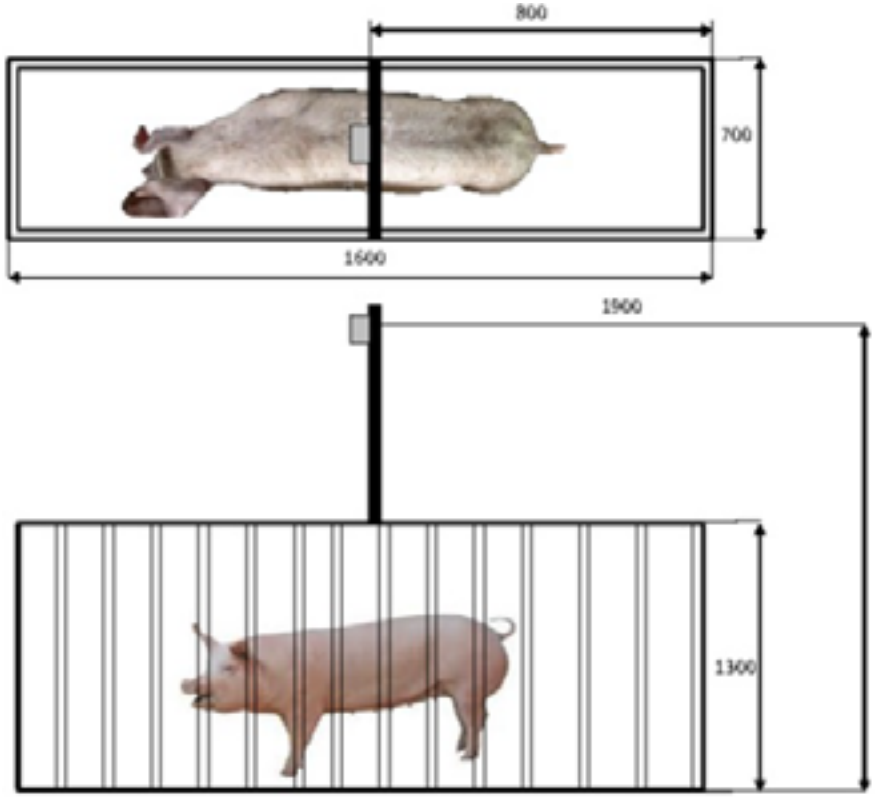

After correcting this information we then could check the frequency by which each pig is weighed independently to get some insights about the reliability of the data. The histogram of Figure 6 shows us the frequency of appearance of each pig. It can be observed that there are only 3 pigs of which have been weighed more than 30 times, which may be an indication that our data may be biased in this sense.

However, these samples correspond to a limited period of time. This may be an explanation for the pigs with few appearances, which may have been heavy just at the beginning or at the end of this period. However it is clear that the samples should be more homogeneous. There is room for improvement in the weighing circuit.
Next, the data was processed in order to exploit its maximum potential, eliminating redundant information and trying to correct all possible errors.

Figure 6. Histogram of the number a pig is weighed

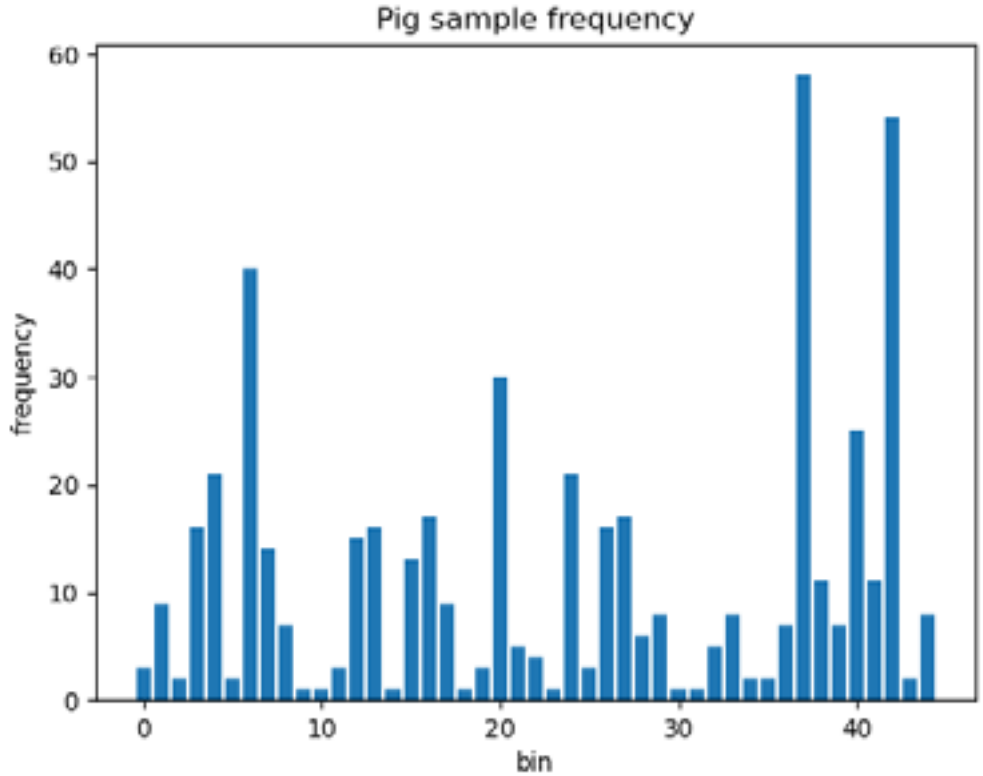

3.1.3 Implementation.

For the first approximation of the weight, a linear regression was setup that consists of different attributes that seem interesting when determining the weight of a pig:

- Area in pixels
- Maximum height
- Mesh area
- Average height
- Approximate volume

Once this model has been trained with the normalized data, we observed the coefficients of the model and they revealed that the maximum height was the attribute with the highest coefficient, that is, the most important, followed by the average height and the number of pixels. This regression approach results in solution that has a Mean Average Error (MAE) of 6.43kg.

Figure 7. Aproximating the shape of a pig with an ellips

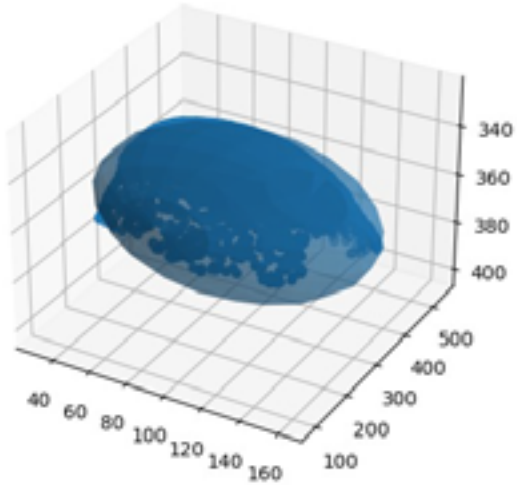

Convolutional Neural Network

Next, we decided to change the lineal regression model of the previous section using a CNN for the regression. The first regressions using CNN were done with the help of the ktrain library. It has built-in early-stopping and also automatic adjustments of the learning rate. First, some initial experiments were done with the raw depth images and while seeking to improve these results, the idea arose to use the segmentation masks on these images to eliminate irrelevant information related to the background.

With the trained models, the images with a higher MAE were reviewed and we came to the conclusion that the variation of the position of heads might be an explanation for some of the variance in the results meaning that the pigs with the head hidden tended to weigh less and conversely the pigs with the head up high weighed more. To apply this correction, the masks of the pigs without the head were used.

Next, we also introduced some data augmentation techniques were also used. First we tried to add some Gaussian noise to the image. The idea was to make model more robust, but it only worsened the result and it was discarded. What did work correctly and was used during development is the use of a vertical flip to augment the data (a horizontal flip made no sense, since all the pigs were oriented the same way).

During development, the idea arose to train a model to classify pigs according to their size in small, medium and large, as a first step in order to later create individual regressive models for each subset (small, medium, large). This classification model obtains an accuracy of 84% and with the regressive models generated after the classification we saw the improvement of the MAE in the groups of small and large pigs. However, in the medium class, as they contain more erroneous classifications, both from the small ones as well as from the large ones, we obtain a much worse MAE, which on average does not improve the strategy of using the set of all the images for the regression, so for that reason this plan was discarded. In Table 1 we give an overview of the results obtained with the different methods using the CNN.

Table 1. Final results

Method	MAE (kg)
Raw images	7.6
Segmentation taking into account the head	4.6
Segmentation removing the head	4.5
Pig centered, taking into account the head	4.2
Pig centered, removing the head	3.6

3.2. Challenges solved and results obtained.

The main conclusion of the project is that we have been able to develop a tool capable of estimating the weight of pigs. We consider that it is a valid solution for the data provided and with a possible future implementation in smart farms. After different attempts and having exploited all the possibilities offered by the data, we have been able to generate a model capable of predicting weight with a MAE of 3.6kg with only 600 samples.

Regarding the treatment of the data, different solutions have been explored when treating the images until finding the combination that gives us the best results. We ended up eliminating the back-

ground which from the first moment we thought could introduce errors, we have centered the pig which facilitates the task for CNN and finally we removed the head in the segmentation of the pig, since its position, by varying so much, ends up introducing an error that worsens our result.

Regarding the regression part, everything indicates that with more samples we would improve the result regardless of the overfitting, due to the fact that all the photos are taken under the same conditions, that is, the same height, the same angle, the same light, for which we would not have problems regarding overtraining, since we are developing the project specifically with the conditions of this farm.

In addition, anomalies in the data, irregularities in the number of samples from each pig and images without weight entry have been detected. This opens the door for us to search for errors in the system that records the data. In the case of errors with respect to obtaining data, it would be another justification for the average error obtained. If this class of problems could be resolved, the solution could be improved.

In this way we closed the project satisfied with the fulfillment of the objectives set at the beginning, despite not having achieved everything we expected. We have complied with obtaining the best possible error trying to take advantage of the data, but dissatisfied with applying the GNN technology for weight estimation, which had to be discarded in the middle of the project due to the lack of information regarding the numerical regression from a graph and the little support of the existing GNNs for its manipulation.

Another option is the use of Long Short-term Memory, networks capable of obtaining predictions taking into account past outputs, in this way we can obtain a prediction more in line with the individual evolution of the animal.

3.3. Current limitations of the Technology.

The current solution is written in Python and is based on Keras. In order to come up with a real-time solution for an industrial environment, the solution has to be ported to a native language. One possible solution is to convert our solution (architecture and weights) to TensorRT. The network has been trained using a specific dataset that corresponds to the pigs of a very particular pig farm. To generalize the solution, pigs of other pig farms should be analyzed as well.

4. Potential impact of the solution

The final idea would be to implement the solution in pig farms to monitor the evolution of the weight of the individual pigs in time. Then the system automatically detects when a pig is ready to be transported to the slaughterhouse. However, the system also helps to detect anomalies in the evolution of the animals (a possible disease) and also the feed for the pig can be adjusted as a function of its weight.

Índex

5. Signatures

Realitzat per:
Coen Antens
Marc Maldonado
Juan Ramón Jiménez
Richard Segovia Barreales

Revisat per:
Coen Antens

Índex

6. Bibliography

AHDB Pork "Pigs 2022 - Pig weighing systems" [Online 30/9/2021]
https://www.youtube.com/watch?v=3CuLNYhFOjk

K. Kollis, C.S. Phang, T.M. Banhazi and S.J. Searle
"Weight estimation using image analysis and statistical modelling: A preliminary study" [Online 30/9/2021]
https://www.researchgate.net/publication/236886740 Weight Estimation Using Image Analysis and Statistical Modelling A Preliminary Study

Dorota Anglart "Automatic estimation of body weight and body condition score in dairy cows using 3D imaging technique" [Online 30/9/2021]
https://stud.epsilon.slu.se/6355/1/anglart d 140114.pdf

Jianlong Zhang, Yanrong Zhuang, Hengyi Ji, and
Guanghui Teng "Pig Weight and Body Size Estimation
Using a Multiple Output Regression Convolutional
Neural Network: A Fast and Fully Automatic Method"
[Online 30/9/2021]
https://www.researchgate.net/publication/352102505 Pig Weight and Body Size Estimation Using a Multiple Output Regression Convolutional Neural Network A Fast and Fully Automatic Method

Eric T. Psota, Mateusz Mittek, Lance Pérez and TyB Schmidt "Multi-Pig Part Detection and Association with a Fully Convolutional Network" [Online 30/9/2021]
https://www.researchgate.net/publication/352102505 Pig Weight and Body Size Estimation Using a Multiple Output Regression Convolutional Neural Network A Fast and Fully Automatic Method

Kaiming He, Xiangyu Zhang, Shaoqing Ren and Jian
Sun "Deep Residual Learning for Image Recognition"
[Online 10/11/2021]
https://arxiv.org/pdf/1512.03385.pdf

Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov and Liang-Chieh Chen
"Mobile-NetV2: Inverted Residuals and Linear Bottlenecks"
[Online 10/11/2021]
https://arxiv.org/pdf/1801.04381.pdf

Olaf Ronneberger, Philipp Fischer and Thomas Brox
U-Net: Convolutional Networks for Biomedical Image
Segmentation" [Online 14/12/2021]
https://arxiv.org/pdf/1505.04597.pdf

CIDAI Centre of Innovation for Data tech and Artificial Intelligence

C/Bilbao, 72 Edif. A. 08005 - Barcelona
Tel. +34 93 7419 100
info@cidai-catalonia-ai.eu
www.cidai.eu

Generalitat de Catalunya
Departament de la Vicepresidència i de Polítiques Digitals i Territori

Ajuntament de Barcelona
Barcelona Activa

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

CVC Computer Vision Center

eurecat Centre Tecnològic de Catalunya

i2cat

Microsoft

NTT DATA

SAP

sdg group

Universitat Politècnica de Catalunya BarcelonaTech
Intelligent Data Science and Artificial Intelligence Research Center

IDEAI