

Tracking Grow-Finish Pigs Across Large Pens Using Multiple Cameras

Aniket Shirke^{a,*}, Aziz Saifuddin^a, Angela Green-Miller^c, Isabella Condotta^b, Achleshwar Luthra^e, Jiangong Li^c, Xiaodan Hu^d, Tawni Williams^b, Aneesh Kotnana^a, Okan Kocabalkanli^a, Narendra Ahuja^d, Ryan N. Dilger^b, Matthew Caesar^a

^a Department of Computer Science, University of Illinois at Urbana-Champaign, Champaign, USA

^b Department of Animal Sciences, University of Illinois at Urbana-Champaign, Champaign, USA

^c Department of Agricultural and Biological Engineering, University of Illinois at Urbana-Champaign, Champaign, USA

^d Department of Electrical and Computer Engineering, University of Illinois at Urbana-Champaign, Champaign, USA

^e Department of Electrical and Electronics Engineering, BITS Pilani, India

* Corresponding author. Email: anikets@illinois.edu

Abstract

Increasing demand for meat products combined with farm labour shortages has resulted in a need to develop new real-time solutions to monitor animals effectively. Significant progress has been made in continuously locating individual pigs using tracking-by-detection methods. However, these methods fail for oblong pens because a single fixed camera does not cover the entire floor at adequate resolution. We address this problem by using multiple cameras, placed such that the visual fields of adjacent cameras overlap, and together they span the entire floor. Avoiding breaks in tracking requires inter-camera handover when a pig crosses from one camera's view into that of an adjacent camera. We identify the adjacent camera and the shared pig location on the floor at the handover time using inter-view homography. Our experiments involve two grow-finish pens, housing 16-17 pigs each, and three RGB cameras. Our algorithm first detects pigs using a deep learning-based object detection model (YOLO) and creates their local tracking IDs using a multi-object tracking algorithm (DeepSORT). We then use inter-camera shared locations to match multiple views and generate a global ID for each pig that holds throughout tracking. To evaluate our approach, we provide five 2-minutes long video sequences with fully annotated global identities. We track pigs in a single camera view with a Multi-Object Tracking Accuracy and Precision of 65.0% and 54.3% respectively and achieve a Camera Handover Accuracy of 74.0%. We open-source our code and annotated dataset at <https://github.com/AIFARMS/multi-camera-pig-tracking>

Keywords: computer vision, continuous, surveillance, precision livestock monitoring, homography.

1. Introduction

Pigs are one of the most raised livestock animals in the world, forming a primary protein source for millions of people across numerous cultures and geographical regions (Shahbandeh, 2020). The increase in demand for high-quality protein has resulted from growing populations combined with increasing incomes across the globe.

Health is a contributor to efficient production, and ideally could be best done by monitoring them individually. Behavior has been well-established as a strong indicator of pig health (Stephen et al., 2016). Manual monitoring of individual behavior in a commercial livestock farm is not practical or sustainable with existing farm staffing and workflow (Swan, 2019). Each animal typically receives no more than a few seconds of observation time on an average day (PIC North America, 2014).

Monitoring pigs automatically requires tracking each animal and identifying and interpreting its behavior. This presents several key challenges. Pens often contain a large number of pigs, that are similar in appearance and thus difficult to visually distinguish from each other. Moreover, pens are often in buildings with low ceilings or otherwise congested areas, making it difficult to see the entire pen from a single vantage point, requiring monitoring across different views, from different vantage points.

To address this challenge, we have developed a system to monitor pigs using multiple cameras with adjacent cameras having overlapping fields of view. We detect and track pigs in each view using state-of-the-art object detection and tracking models. We then use homography between adjacent camera views to identify a pig in their overlapping views, transitively match different camera views of the same pig, and thus assign a global identity to each pig, thereby achieving global tracking. In order to validate our methods, we have monitored two grow-finish pens using a set of appropriately chosen and placed cameras.

Our primary contributions in this paper can be summarized as follows:

- We present a multi-camera pig tracking system for identifying each pig across all camera views. We are not aware of any other such system.
- We present a multi-camera, multi-pen dataset containing videos of pigs in multiple pens captured from multiple viewpoints and make it publicly available. We are not aware of any such publicly available dataset.

2. Related Work

Individual tracking of animals in group housing is a demanding task. Uninterrupted tracking poses significant challenges when there is a lack of discernible differences in the physical characteristics of the animals. To address these challenges, tracking-by-detection methods have been proposed. Psota et al. (2014) cast detection as a segmentation task. The four semantic parts of pig (ears, shoulder, and tail) are detected and tracked using a Fully Convolutional Network and the Hungarian algorithm. In addition, the authors provide a dataset with 2000 images from multiple pens, which is publicly available.

In Seo et al. (2020), a TinyYOLO (Redmon et al., 2016) architecture is employed to detect pigs from infrared videos, with emphasis on execution speed as the target platform is an embedded device. Jaewon et al. (2019) propose a method to detect pigs under various illumination conditions, by combining information from the depth and infrared images using spatio-temporal interpolation. Similarly, in Johannes et al. (2020) the bounding boxes are replaced with ellipses, which are detected through a segmentation network. Zhang et al. (2019) use an SSD (Wei et al., 2016) architecture coupled with the MOSSE (Bolme et al., 2010) algorithm to perform animal tracking.

Prior work in multi-camera, multi-target tracking typically assumes that the target has distinguishing features, which can be used for re-identification across multiple views (Cabrera et al., 2011; Ristani and Tomasi, 2018). But this has not led to a similar ability to track pigs because of the lack of such visual features on pigs. To the best of our knowledge, our work is the first application of multi-camera tracking to pigs. Moreover, the only few publicly available datasets for pig monitoring consist of a single-camera view (Psota et al. 2014; Bergamini et al. 2021; Riekert et al. 2020), with no open-source implementations. To address this need, we have collected and open-sourced our multi-camera view dataset, along with our implementation.

3. Methods

3.1. Detection and Local Tracking

For each camera view, pig detection is achieved by using the state-of-the-art YOLOv4 (Bochkovskiy et al., 2020) model. In order to generate tracking IDs in a single camera view, the detections provided by YOLOv4 are tracked using DeepSORT (Wojke et al., 2017). The following subsections describe how local tracking IDs are processed to obtain global tracking IDs for pigs.

3.2. Homography Estimation

Any two images of the same planar surface are related by a homography. In our case, the planar surface is the pen floor, observed from a *Ceiling* view and an *Angled* view, as depicted in Figures 1a and 1b, respectively.

Assuming that the *Angled* view is aligned with the world coordinate axes and has no translation, a 3-D point P is mapped to the 2-D pixel point p_{angled} in the *Angled* view using the following relation:

$$\lambda_1 * p_{angled} = K [I | 0] P$$

where K is the intrinsic parameter matrix, I is a 3x3 identity matrix, and p_{angled} and P are homogeneous coordinates defined as follows:

$$p_{image} = \begin{bmatrix} x_{image} \\ y_{image} \\ 1 \end{bmatrix} \quad P = \begin{bmatrix} X_{world} \\ Y_{world} \\ Z_{world} \\ 1 \end{bmatrix}$$

Similarly, the same 3-D point P is mapped to the 2-D pixel point $p_{ceiling}$ in the *Ceiling* view using the relation:

$$\lambda_2 * p_{ceiling} = K [R | t] P$$

where R and t is the relative rotation and translation of the *Ceiling* view with respect to the *Angled* view. Note that λ_1 and λ_2 are free parameters as multiple 3-D points can get mapped to the same 2-D point in the image due to single-view ambiguity.

Our task is to estimate a homography $H(\cdot)$ such that

$$p_{angled} = H(p_{ceiling})$$

In an ideal scenario, having many 3D-2D correspondences between the floor pen and the corresponding image views can help in obtaining the homography accurately. Obtaining those correspondences is challenging at a farm due to the large numbers of pigs occluding the floor. In the absence of such correspondences, multiple methods can be potentially used to estimate homography between two camera views.



Figure 1: Simplified depiction of the Multi-Camera Tracking algorithm

We adopt the following method to estimate the homography $H_{ceiling \rightarrow angled}$:

1. Using a perspective transformation, the *Ceiling* view is first transformed into a top-down view parallel to the pen floor. The homography between these two views is denoted as $H_{ceiling \rightarrow top-ceiling}(\cdot)$
2. The *Angled* view is then transformed into a top-down view using a similar perspective transformation. The homography between these two views is denoted as $H_{angled \rightarrow top-angled}(\cdot)$
3. Since both the top-down views are parallel to the pen floor, key points from the overlapping regions of both these views are matched. These key points can be detected as well as matched automatically. But in this paper, we perform these steps manually, by inspection, and estimate a homography using RANSAC (Derpanis, 2010). The homography is denoted as $H_{top-ceiling \rightarrow top-angled}(\cdot)$

Thus, a pixel point $p_{ceiling}$ in the *Ceiling* view can be mapped to a point $p_{estimated}$ in the *Angled* view as follows:

$$\begin{aligned} p_{estimated} &= H_{angled \rightarrow top-angled}^{-1}(H_{top-ceiling \rightarrow top-angled}(H_{ceiling \rightarrow top-ceiling}(p_{ceiling}))) \\ &= H_{ceiling \rightarrow angled}(p_{ceiling}) \end{aligned}$$

3.3. Track Aligning Algorithm

Algorithm 1: Multi-Camera Tracking algorithm

Data: Set of c and a local pig IDs and bounding boxes in *Ceiling* view and *Angled* view respectively:
ceiling_tracks, angled_tracks

Result: Dictionary of matches between ceiling_ids and angled_ids: matches

/ Generate transformed_ceiling_tracks by using the homography $H_{ceiling \rightarrow angled}$ */*

matrix = [] //size of the matrix will be $c \times a$

For (ceiling_id, bbox) **in** transformed_ceiling_tracks:
 matrix.append(intersection(bbox, angled_tracks))

/ Greedily pop the max value from the matrix and enter in matches */*

matches = {}

While non-zero entry in matrix exists:

 ceiling_id, angled_id = matrix.argmax()

 matches[:, angled_id] = 0

If ceiling_id not in matches:

 matches[ceiling_id] = angled_id

We use the homography $H_{ceiling \rightarrow angled}$ computed in the previous subsection to develop a multi-camera tracking algorithm, by aligning tracks in the *Ceiling* view and the *Angled* view.

The bounding box from the *Ceiling* view can be treated as a quadrilateral on the pen floor in which the corresponding pig is contained. The bounding box from the *Ceiling* view is then projected under the homography to a quadrilateral in the *Angled* view, as depicted in Figure 1c. The transformed quadrilateral is then matched with a bounding box in the *Angled* view with which it has significant overlap in terms of the pixel area. As seen in Figure 1b, ID 56 in the *Ceiling* view will be matched with ID 6 in the *Angled* view.

4. Evaluation

4.1. Camera Deployment




Figure 2: Camera Deployment at the Imported Swine Research Laboratory, UIUC



Figure 3: (Left) Angled view of Pen C. (Right) Ceiling view of Pens B and C.

Videos were recorded in two grow-finish pens (B and C) at the University of Illinois, Urbana-Champaign’s Imported Swine Research Laboratory (ISRL), which serves as the research testbed for our experiments. Our surveillance setup monitored 17 and 16 pigs raised in pens B and C respectively. Three wide-angle lens cameras were strategically placed to cover the drinkers and feeders in the visual field, as shown in Figure 2. The *Ceiling* camera covered the drinkers and was placed at a height of 4 metres. The *Angled* cameras covered the feeders and were placed at a height of 2.2 metres. Figure 3 depicts the *Angled* and *Ceiling* camera views. Infrared floodlights were deployed to enhance the existing night vision capabilities of the cameras. The video feed was captured 24x7 at 4k resolution and 15 frames per second.

4.2. Dataset

For training the YOLOv4 model, we sample and annotate 429 images using the VGG Image Annotator (Dutta et al., 2016). The images were randomly split in an 80:20 between the training set (343 images) and the validation set (86 images). To evaluate the tracking efficacy, five 2-minutes long video sequences were annotated with global pig IDs using a custom MATLAB tool. The annotations for two camera views are made every 15th frame, thus (5*2*60*2) 1200 annotated frames are available, in addition to the 429 images for object detection.

4.3. YOLOv4 detection

We train the YOLOv4 model on an Nvidia V100 GPU provided by the HAL computing cluster (Kindratenko et al., 2020). The model was configured for an input resolution of 608x608 and optimized for 2000 iterations with a batch size of 64, a learning rate of 0.001, and momentum of 0.95. We evaluate the model using standard object detection metrics. A mean Average Precision of 99.5% and an average Intersection over Union of 80.52% is achieved on the validation set.

4.4. Tracking evaluation

We evaluate the efficacy of the single-camera tracking algorithm (DeepSORT) using standard Multi-Object Tracking Metrics (Ristani et al., 2016). The Multi-Object Tracking Accuracy (MOTA) takes into account three sources of errors and can be defined by the following equation: $MOTA = 1 - [\sum_t (FN_t + FP_t + IDSW_t) / \sum_t GT_t]$ where FN_t , FP_t and $IDSW_t$ are False Negatives, False Positives and Identity Switches, and GT_t is the ground truth number of bounding boxes at time t . Multi-Object Tracking Precision (MOTP) captures the localization precision of the object detector. Standard

metrics such as Precision and Recall measure the number of mismatched or unmatched detection-frames, regardless of where the discrepancies start, or end or which cameras are involved. We also report identity related metrics computed after global min-cost matching: IDF1, IDP, and IDR. IDF1 is the ratio of correctly identified detections over the average number of ground-truth and computed detections, IDP (precision) is the fraction of computed detections that are correct whereas IDR (recall) accounts for correctly identified ground truth detections.

Type	Local Tracking (DeepSORT)							Global Tracking
	IDF1	IDP	IDR	Recall	Precision	MOTA	MOTP	CHA
Day	66.1	61.8	71.0	98.0	85.4	80.6%	61.3%	92.4%
Night	53.2	49.9	56.9	85.6	75.1	55.2%	44.5%	46.3%
Overall	58.2	54.5	62.3	90.4	79.0	65.0%	54.3%	74.0%

Table 1: Local and Global Tracking metrics

We evaluate our multi-camera tracking approach by reporting the Camera Handover Accuracy (CHA). CHA is defined as the fraction of predicted identity matches for two camera views out of the total ground truth matches. Intuitively, MOTA and MOTP quantifies how accurately the algorithm is tracking pigs with minimal false positives, and CHA quantifies how accurately identities are exchanged between cameras using homography. As seen in Table 1, we observe that tracking is relatively easier during day as the quality of video feed degrades at night. We achieve an MOTA, MOTP and CHA of 65%, 54.3% and 74% respectively.

5. Discussion

5.1. Issues in Detecting Pigs using YOLOv4




Figure 4: False Positives and False Negatives in Detection

Occlusions are a primary source of false negatives. Even though detecting pigs by a direct application of state-of-the-art deep learning models, such as YOLOv4, yield great results, there are several scenarios where detection of pigs becomes very challenging. One of the most common impediments to detection is occlusion. If a pig is partially or completely occluded in any of the views, then the model tends to either miss the pig completely or predicts a bigger bounding box, encompassing more than one pig, which leads to missed detections. Having two camera views helps in partially alleviating this problem.

False positives result from part of pig being recognized as a whole pig. Pigs tend to huddle in pens to keep themselves warm as a group, which creates a challenge for recognition. Usually, a post-processing non-maximal suppression step during object detection helps in selecting the best bounding box with the maximum confidence score. But often, the model considers a part of the pig to be the whole pig and predicts a box that is not suppressed during non-maximal suppression.

One can see both kinds of errors, as described above, in Figure 4. The pig in the bottom left of the frame is not detected by the model due to heavy occlusion from other pigs and the pen door. One way to mitigate this issue can be to augment the dataset by emulating the artifacts encountered in the pen environment, such as occlusions from grid-like doors in this case. Additionally, detections ‘Pig-22’ and ‘Pig-15’ in Figure 1 belong to the same pig. Here, ‘Pig-22’ predicts a bounding box for a part of a pig and should be ideally suppressed. It is also worth noting that ‘Pig-16’ is a false detection as the two hind legs of two different pigs confuse the model into believing that it is a single pig.

5.2. Issues in Tracking Pigs using DeepSORT




Figure 5: Identity Swap in DeepSORT

The detection issues described in the previous subsection affect the tracking performance of DeepSORT, as false positives lead to an extra number of tracks and false negatives lead to premature termination of tracks. Another primary source of tracking errors is induced by rapid movements in a pig huddle. As seen in Figure 5, the identity of Pigs 11 and 16 get swapped due to a rapid social interaction with Pig 12. Theoretically, DeepSORT should tackle and minimize such identity swaps as it uses an appearance model for re-identification. But the lack of visually distinguishing features between pigs calls for advanced re-identification algorithms while tracking. One potential solution might be to apply another recognition approach, such as a gait analysis, to reidentify them upon occlusion.

5.3. Issues in Global Tracking

The inherent problems in detection and tracking described in the previous subsections affect global tracking, as false detections can lead to false correspondences between different views. The algorithm is also sensitive to variations in detections from a pig huddle, as any error in homography estimation can magnify and lead to an identity swap between multiple pigs. This issue can be mitigated by precise homography estimation and by only assigning a global ID for pigs that have higher overlapping confidence. But that can lead to missed opportunities of matching, leading to a low recall value. Additionally, the multi-camera algorithm described in section 3.3 works under the assumption that all the cameras are synchronized with each other. But in actual deployments, there is a variable delay of several minutes between the capture times of the Ceiling and the Angled views. The delay needs to be fixed manually as such homography-based multi-camera trackers need tight synchronization between the two views.

6. Conclusions

Livestock monitoring is becoming increasingly important in precision livestock management. In order to tackle the problem of monitoring livestock residing in large pens, we propose a scheme for homography-based multi-camera tracking. We achieve a single-camera tracking accuracy and precision of 66.7% and 76.2% respectively, along with a camera handover accuracy of 74%. We plan to further enrich our multi-camera tracking dataset with key focal points and a validated ethogram for pig behavior. We open-source this dataset so that it can be leveraged to build educational applications for pig tracking and behavior monitoring.

References

- [1] PIC North America. Standard animal care: Daily routines. Wean to Finish Manual, pages 23–24, 2014. 1
- [2] Luca Bergamini, Stefano Pini, Alessandro Simoni, Roberto Vezzani, Simone Calderara, Rick BD Eath, and Robert B Fisher. Extracting accurate long-term behavior changes from a large pig dataset. In 16th International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications, VISIGRAPP 2021, pages 524–533. SciTePress, 2021.
- [3] Alexey Bochkovskiy, Chien-Yao Wang, and Hong-Yuan Mark Liao. Yolov4: Optimal speed and accuracy of object detection, 2020.
- [4] David S Bolme, J Ross Beveridge, Bruce A Draper, and Yui Man Lui. Visual object tracking using adaptive correlation filters. In 2010 IEEE computer society conference on computer vision and pattern recognition, pages 2544–2550.

IEEE, 2010. 2

- [5] Johannes Brünger, Maria Gentz, Imke Traulsen, and Reinhard Koch. Panoptic instance segmentation on pigs. arXiv preprint arXiv:2005.10499, 2020.
- [6] Reyes Rios Cabrera, Tinne Tuytelaars, and Luc Van Gool. Efficient multi-camera detection, tracking, and identification using a shared set of haar-features. In CVPR 2011, pages 65–71, 2011.
- [7] A. Dutta, A. Gupta, and A. Zissermann. VGG image annotator (VIA). <http://www.robots.ox.ac.uk/vgg/software/via/>, 2016. Version: X.Y.Z, Accessed: 04/14/2021.
- [8] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott Reed, Cheng-Yang Fu, and Alexander C. Berg. Ssd: Single shot multibox detector. Lecture Notes in Computer Science, page 21–37, 2016.
- [9] Stephen G. Matthews, Amy L. Miller, James Clapp, Thomas Plötz, and Ilias Kyriazakis. Early detection of health and welfare compromises through automated detection of behavioural changes in pigs. The Veterinary Journal, 217:43–51, 2016.
- [10] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi. You only look once: Unified, real-time object detection. In 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pages 779–788, 2016.
- [11] Martin Riekert, Achim Klein, Felix Adrion, Christa Hoffmann, and Eva Gallmann. Automatically detecting pig position and posture by 2d camera imaging and deep learning. Computers and Electronics in Agriculture, 174:105391, 2020.
- [12] Ergys Ristani, Francesco Solera, Roger S. Zou, Rita Cucchiara, and Carlo Tomasi. Performance measures and a data set for multi-target, multi-camera tracking, 2016.
- [13] Ergys Ristani and Carlo Tomasi. Features for multi-target multi-camera tracking and re-identification. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), June 2018.
- [14] Jaewon Sa, Yunchang Choi, Hanhaesol Lee, Yongwha Chung, Daihee Park, and Jinho Cho. Fast pig detection with a top-view camera under various illumination conditions. Symmetry, 11(2):266, 2019.
- [15] Jihyun Seo, Hanse Ahn, Daewon Kim, Sungju Lee, Yongwha Chung, and Daihee Park. Embeddedpigdet—fast and accurate pig detection for embedded board implementations. Applied Sciences, 10(8):2878, 2020.
- [16] M Shahbandeh. Number of pigs worldwide from 2012 to 2020. <https://www.statista.com/statistics/263963/number-of-pigs-worldwide-since-1990>.
- [17] M. K. Swan. Swine human resources: Managing employees. <https://swine.extension.org/swine-human-resources-managing-employees>.
- [18] Eric T Psota, Ty Schmidt, Benny Mote, and Lance C Pérez. Long-term tracking of group-housed livestock using key-point detection and map estimation for individual animal identification. Sensors, 20(13):3670, 2020.
- [19] N. Wojke, A. Bewley, and D. Paulus. Simple online and realtime tracking with a deep association metric. In 2017 IEEE International Conference on Image Processing (ICIP), pages 3645–3649, 2017.
- [20] Lei Zhang, Helen Gray, Xujiang Ye, Lisa Collins, and Nigel Allinson. Automatic individual pig detection and tracking in pig farms. Sensors, 19(5):1188, 2019.
- [21] Volodymyr Kindratenko, Dawei Mu, Yan Zhan, John Maloney, Sayed Hadi Hashemi, Benjamin Rabe, Ke Xu, Roy Campbell, Jian Peng, and William Gropp. Hal: Computer system for scalable deep learning. In Practice and Experience in Advanced Research Computing, PEARC '20, page41–48, New York, NY, USA, 2020. Association for Computing Machinery
- [22] Derpanis, Konstantinos G. "Overview of the RANSAC Algorithm." Image Rochester NY 4, no. 1 (2010): 2-3.