A data infrastructure to collect and store raw sensor data (tracking and video footage) to continuously monitor (behaviour related to) locomotion and Digital Dermatitis in dairy cattle

The dot on the horizon of this research was to enable the monitoring of a large group of animals for a long period of time using either tracking data and/or video footage with minimum input from people and having minimal invasive impact on the animals under study. Locomotion of dairy cattle was selected as use-case; lameness or abnormal behaviour due to lameness have a negative impact on animal health and welfare, and production. Monitoring lameness is often done through visual assessment of abnormal locomotion, but visual assessment is subjective and time consuming. Also, very mild cases of lameness and subtle changes in locomotion are difficult to see. The objective of this study was two-fold: (1) develop an infrastructure at our research facility Dairy Campus (www.dairycampus.nl) that allows continuous collection, storage, and synchronisation of high-quality raw sensor data, (2) develop tools to retrieve relevant information for the continuous monitoring of (behaviour related to) locomotion from these raw sensor data.

Technology used in this research included an ultrawideband tracking device and cameras. The tracking device (Tracklab, Noldus Information Technology [1]) was installed such that raw XY-position data of 112 animals were recorded 24/7. Cameras were installed at three different locations at Dairy Campus: (1) one research unit housing 16 cows was equipped with eight Axis cameras (Noldus Information Technology [2]) offering bird-view video footage (from 5am to 6pm), (2) two Hikvision IP cameras [3] were installed in the exit of the milking parlour, allowing the collection of side-view video footage of the same 112 cows that were fitted with a tracking device when walking back from the milking parlour to the barn twice daily, for three hours per milk session, and (3) two Reolink IP cameras in the milking parlour itself [4], allowing the collection of video-footage of all ~450 dairy cows milked in the rotary, while standing on the rotary platform twice a day.

A data architecture was developed allowing raw sensor data collection and data quality control pipelines at different locations and at different levels. For example, to ensure proper functioning of hardware (e.g., ultrawideband anchors and cameras), a ping was requested every minute from every device. If no pings were generated, an email was sent to the data engineer indicating which hardware was not functioning. Another level of data quality control involved the control of tracking tags. A digital tool was developed to record which tag was fitted to which cow, that showed the battery status of the tag, and that checked whether cows went back to the research unit they were supposed to be in when returning from milking. Another level of data quality control involved the data collection and storage pipelines. Tools for monitoring and controlling data quality are detailed in [5]. Tools were developed that helped going from (1) raw tracking data to one XY-position per cow per second, and from that to 24h activity bouts per cow with a tracking device [6], (2) from raw video footage to gait features that can be used for cow individual locomotion monitoring over time [7], (3) raw video footage to the tracking of 3D key points of individual cows housed in a group, allowing the assessment of behaviour (e.g., time to get up or lie down [8], (4) from raw video footage to classification of DD status [9], and (5) from using the processed raw position data to walking trajectories essential for understanding transmission [9].

Tracking devices

Tracklab (Noldus)

All cows in seven barn units (with 16 cows each) in the environmental barn of Dairy Campus were equipped with a Sewio Leonardo personal UWB sensor on their neck collar, placed in tailor made adaptors from Noldus Information Technology (https://www.noldus.com/tracklab). Additional tags and chargers are available to replace and reload tags with (nearly) flat batteries. The barn units and the waiting area next to the milking parlour were equipped with Sewio UWB anchors to communicate with the tags. These anchors were connected to the local fiber network, a dedicated server (PC) collected the messages from the anchors and calculated raw XY-positions from these. The raw position data were stored in a Sewio database, but also copied to a Tracklab server. The UWB system was installed and tuned by Noldus Information Technology personnel.

There were two experiments at WBVR with the ultra-wideband tracking device in combination with the Tracklab software from Noldus Information Technology. One experiment with 17 sheep (tracked for 8 days) and one with 16 pigs (tracked for up to 31 days). Measurement frequency was set at 1Hz.

Camera’s

Axis camera

A total of eight Axis P1375-E cameras (https://www.noldus.com/viso) were installed in one of the research units housing 16 dairy cows at Dairy Campus. The resolution of the cameras was 1920x1080 pixels per frame, recording was set at 25 fps. All cameras were connected to the local fiber network and automatically synchronized with an NTP-server. Video footage was captured using a tailor made ffmpeg-script and temporarily stored on a local server. The cameras were installed such that the home environment of the cows was viewed from 4 different viewpoints in stereo, covering the entire barn lay-out. The cameras were calibrated after initial installation using a 600*800 mm checkerboard with 18*25 squares of 30 mm2 (Calib.io). For determination of the camera extrinsics patterns in the slatted floor elements with known geometry are recorded, such that combinations of pixel coordinates and floor coordinates can be determined from the video footage recorded.

Hikvison camera

Two HikVision IP cameras of model DS-2CD2T45FWD-I5 were installed besides the exit lane where the cows walk through after milking to return to their barn units. These cameras were also automatically synchronized with an NTP-server and connected to the local fiber network. The resolution of these cameras was 2688 *1520 pixels, recording was set at 25 fps. Video footage was captured using a tailor made ffmpeg-script and temporarily stored on a local server.

Reolink cameras

Two Reolink IP cameras of model RLC-811A (www.reolink.com) were installed in a protection device box designed and developed by Eric Karruppannan from Wageningen Technical Solutions group of WUR. The protection box was placed at the end of the milking carrousel at DC milking parlour. These cameras were also automatically synchronized with an NTP-server and connected to the local fiber network and calibrated as described before. The resolution of these cameras was 4K 8MP (3840 x 2160 pixels). Video footage was captured using an open-source FFMPEG software stored on a local server (at DC). A special box container was designed and built as the environment where the cameras had to function was almost comparable to aquatic cameras, in terms of moisture sealing and water protection. This complementary hardware resulted from the experience gained during the project as another set of camaras and connectors were quickly spoiled and no longer working. Another issue was the faeces and urine from the cows causing the cameras to get extremely dirty, quickly. Therefore, so the special box included a transparent window with two purposes: protection of the camera’s from moist, water, and dirt, and facilitating cleaning.

Monitoring locomotion and digital dermatitis in dairy cattle

Data quality monitoring and control

Data quality can be assessed from various perspectives. Firstly, data completeness is crucial. Completeness is determined by the underlying sensors, batteries, and beacons, as well as the processing computers and cameras that collect data from these sensors and beacons. To monitor quality, we've implemented monitoring at different levels. Initially, we monitor at the network level to ensure the underlying network infrastructure is functioning. Every hour, a network ping is sent to all components. If any issues arise, the research coordinators are notified through email. Additionally, we monitor the pipelines (to collect and store raw sensor data in a structured logical manner) to ensure they are running smoothly, and no data is lost. Using Apache Airflow, we gain insight into the pipelines, and if a pipeline fails, alerts are triggered. Next, we verify if the data is reaching the secured W-drive as intended. These processes run continuously and automatically every day. Lastly, for both video and sensor data, manual quality checks have been implemented to assess the content and quality of the data in detail.

Tools from raw position data (X, Y) to behaviour bouts to be used for behaviour monitoring

The data pipeline developed to extract raw XY position data from Sewio, utilizes a series of scripts executed on a regular basis. The MySQL database of Sewio stores both position and accelerometer data, but the database is not optimized for additional querying. To prevent stressing the production systems, a mirrored database has been set up with indexed tables, allowing data to be stored on the W-drive. To extract the data, a script called sync.py is invoked every hour. This script retrieves the last million records from the position and history tables of the MySQL database and saves them to the mirror database. Subsequently, another script called export_to_parquet.py is used to process and convert these data into Parquet files. These Parquet files contain both history and position data and are generated using the fastparquet engine. The choice of Parquet as the final data format was made due to its efficient compression and support for parallel processing, making it easier to store and process large amounts of data.

To simplify the setup and management process of the pipeline, Apache Airflow has been implemented, providing more insight into the processes and enabling effective building and monitoring of pipelines. This allows the pipeline to be scaled and managed with more transparency and control. To utilize this toolbox, data scientists should have basic knowledge of Linux, MySQL, SSH.

After the data pipeline is implemented the raw position data is stored in parquet-files (one file per day). In these parquet-files the tag identification is not linked to the cow identification. This is done in another pipeline that processes the raw position data into meaningful behavioural information. For this, the pipeline uses the recordings of tag/cow combinations with their start and end date/times. It takes into account the time difference between UTF time and CET time. Locations are used to classify behaviours, the following are distinguished for the barn units with cubicles: feeding rack, drinking trough, cubicles A, B and C, concentrate feeder, and walking (see Figure 1). For the unit without cubicles the following areas are distinguished: feeding rack, drinking trough, concentrate feeder, slatted floor and resting area. For all cows the waiting area (an assemble area where cows wait in groups just before getting milked) is an additional location that is identified. From the time series of location records, bouts of behaviour are calculated when duration of the same behaviour fulfils a minimum threshold (see Figure 1). Also, a minimum interval length between bouts is implemented to consider bouts of the same behaviour as separate bouts. In addition to this, activity is calculated as the distance travelled, as the percentage of the day being active, and the percentage of the day a cow is not in the cubicles.

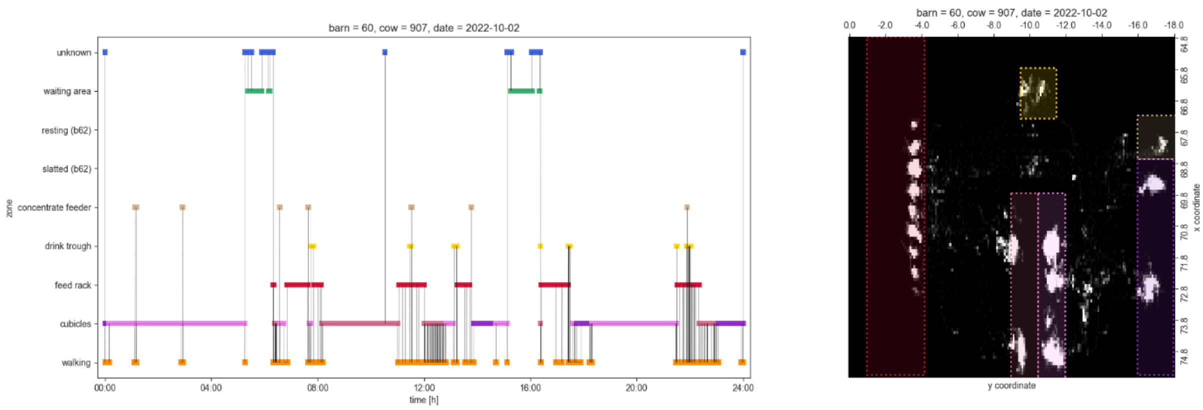

Figure 1. Transforming raw XY positions to real-world locations in the barn (right), and transform these locations into behavioural bouts (left) per cow per day

From raw video footage to gait features to be used to locomotion monitoring

Several steps and scripts have been developed to create a tool for continuous locomotion monitoring. A side-view footage was recorded with 25 frames per second of cows exiting the milking parlor. We created a pipeline for processing raw video footage into gait features. This pipeline includes various steps: (1) cow detection and (2) subsequent keypoint detection using deep learning-based algorithms (YOLO and T-leap), post-processing to correct outliers (3), and (4) extraction of gait features from these post-processed keypoints. More detailed information can be found in (Taghavi et al., 2023). Figure 2 demonstrates the pipeline schematically. In the pipeline, each step takes the output of the previous step as its input. For deep learning models, we utilized a subset of the dataset to fine-tune models tailored to our dataset. In step (1) - cow detection: We trained and tested an object detection deep learning algorithm (YOLOv8) on 200 frames to determine the optimal model for detecting cows passing by the camera. In the subsequent step, keypoint detection: We employed 758 samples of two successive frames to train and test a previously developed deep learning model known as T-LEAP. This model was trained to detect 17 key points on cows within our indoor farm environment, accounting for natural occlusions. In the next step, several postprocessing techniques were applied to correct camera distortion and rectify erroneous key point detections. From the post-processed key points, several gait features were extracted, including stride length, stride duration, and ground duration. By monitoring these gait features for each individual cow, sudden changes in locomotion can be identified and detected.

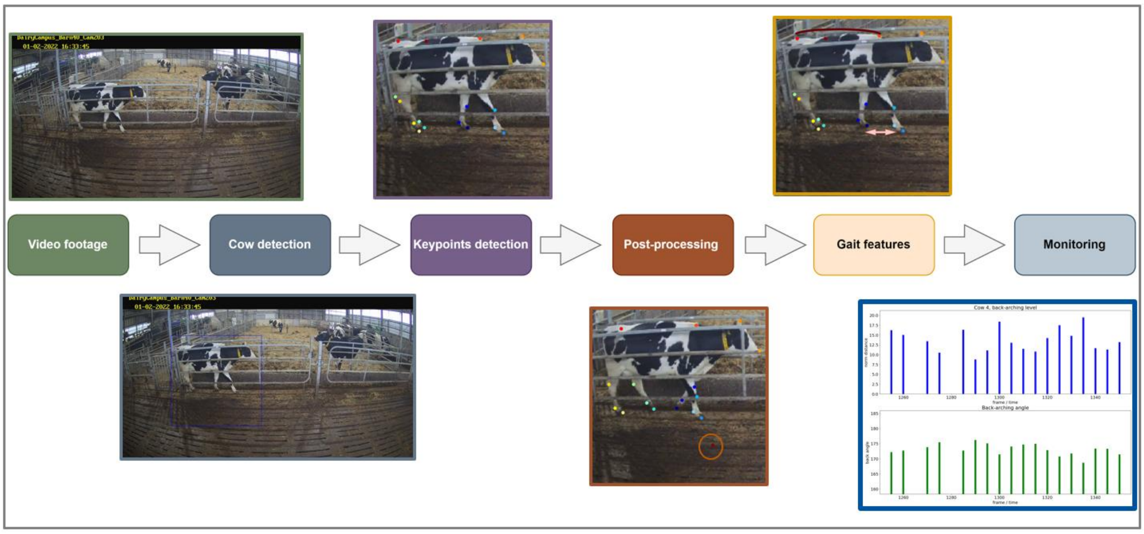

Figure 2. Schematic pipeline developed for locomotion monitoring

From raw video footage to tracked 3D key points to be used for behaviour monitoring

Several steps and scripts have been developed to create a tool to obtain time series of 3D keypoints for behaviour monitoring. Video footage was recorded in one barn unit with 25 frames per second from 4 distinct viewpoints with two cameras for each viewpoint. The cameras per pair were positioned approximately 50 cm next to each other. All cameras were calibrated using the Matlab® calibrator app, camera extrinsic (position and orientation of the cameras) were determined using the pattern of the slatted floor in the barn floor and the camera intrinsics. We created a pipeline for retrieval of 3D keypoints from synchronised raw video footage. This pipeline (see Figure 3) includes various steps: (1) cow detection and (2) subsequent keypoint detection using deep learning-based algorithms (YOLO V8 and MMPOSE respectively), post-processing to determine the combinations of cows in overlapping footage and calculate 3D coordinates of the keypoints (3). The last step of the pipeline (4) is work in progress, and falls under NLAS Toolbox II. For now, we have investigated the usage of time series of keypoint positions to monitor head movements at the feeding rack and are working on monitoring getting up and lying down movements based on keypoint position time series.

In the pipeline, each step takes the output of the previous step as its input. Step 3 also uses the camera intrinsic parameters determined from calibration and the camera extrinsic. The latter have to be checked and if needed be updated when using footage for which the extrinsics were not determined, for both tasks Matlab-scripts are available. For deep learning models, we utilized a subset of the dataset to fine-tune models tailored to our dataset setup. In step (1) - cow detection: We trained and tested an object detection deep learning algorithm (YOLOv8) on 400 frames (390 for training and 10 for validation) to determine the optimal model for detecting cows. For each camera 50 frames were annotated for this task. For step 2 keypoint detection was trained and tested in two iterations. We selected 4 of the 20 keypoints that were annotated for the Animal Pose dataset (https://sites.google.com/view/animal-pose/): right and left ear base, withers and tailbase. In the first iteration 200 frames were annotated (bounding boxes were pre-annotated with the cow-detection model) for the 4 cameras with bird eye view (50 for each camera). Of these, 190 were used for training and 10 for validation. In the second iteration a next set of 120 frames (30 for each camera with bird eye view) was annotated (after pre-annotation with the cow detector and the keypoint detector) with the CVAT annotation tool. The overall set of annotations was checked for consistency with Matlab scripts before they were used for retraining, and inconsistencies detected (e.g. in labelling keypoints as valid, invisible or occluded or as left instead of right ear base and vice versa) were corrected. Hereafter the model was retrained with 300 images and 20 images were used for validation. The final model was applied on several selections of synchronised footage, and procedures to smooth keypoints projections before combining information from two cameras of a pair and for calculation of 3D coordinates from 2 or more keypoint projections are developed. Current work focuses on pattern detection in keypoint time series, with the aim to determine features to describe lying down and getting up events.

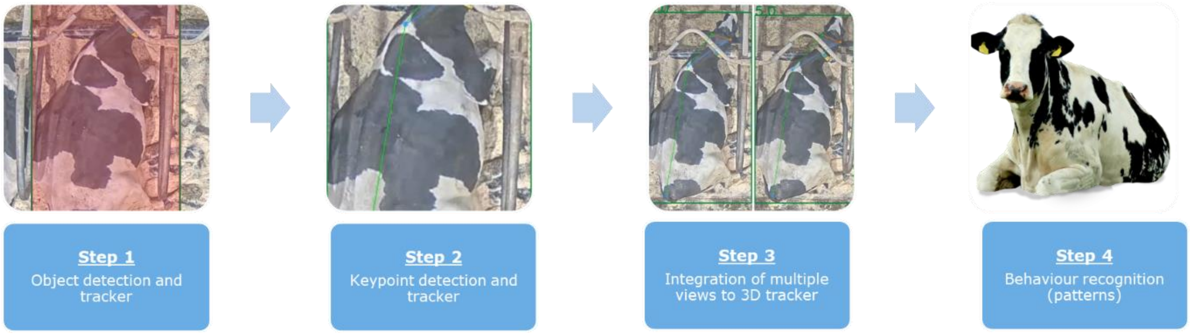

Figure 3. Pipeline for creating timeseries of 3D keypoints.

From raw video footage to classification of the Digital Dermatitis status

Several steps and scripts were developed to create a tool for assessing DD status based on images when evaluating video footages that were recorded during milking. Side-view footage was captured at 25 frames per second of cows as they were about to leave the milking parlour. We created a pipeline to process the raw video footage. This pipeline consists of two steps: (1) cow detection and identification (Figure 4) and (2) subsequent key point detection using deep learning-based algorithms (YOLO V.5). In the pipeline, each step takes the output of the previous step as input. For the deep learning models, we used a set of videos obtained and stored in our dataset from different milking sessions. For step (1) - cow detection and identification, we trained and tested an object detection deep learning algorithm (YOLOv5) on 1200 frames to determine the optimal model for detecting de ID of the carousel station the cow is in when passing by the camera as a proxy for cow ID. This information was subsequently crossed-referenced with the milk recording file from Dairy Campus reporting which cow was present in what carrousel station at what time to retrieve actual cow identification. In the subsequent step, (2) - It was decided whether the quality of the image was considered good enough for trying to get a classification of DD status or not. In this step, we employed 524 claw-samples to train and test a developed deep learning model on DD lesions or a healthy claw.

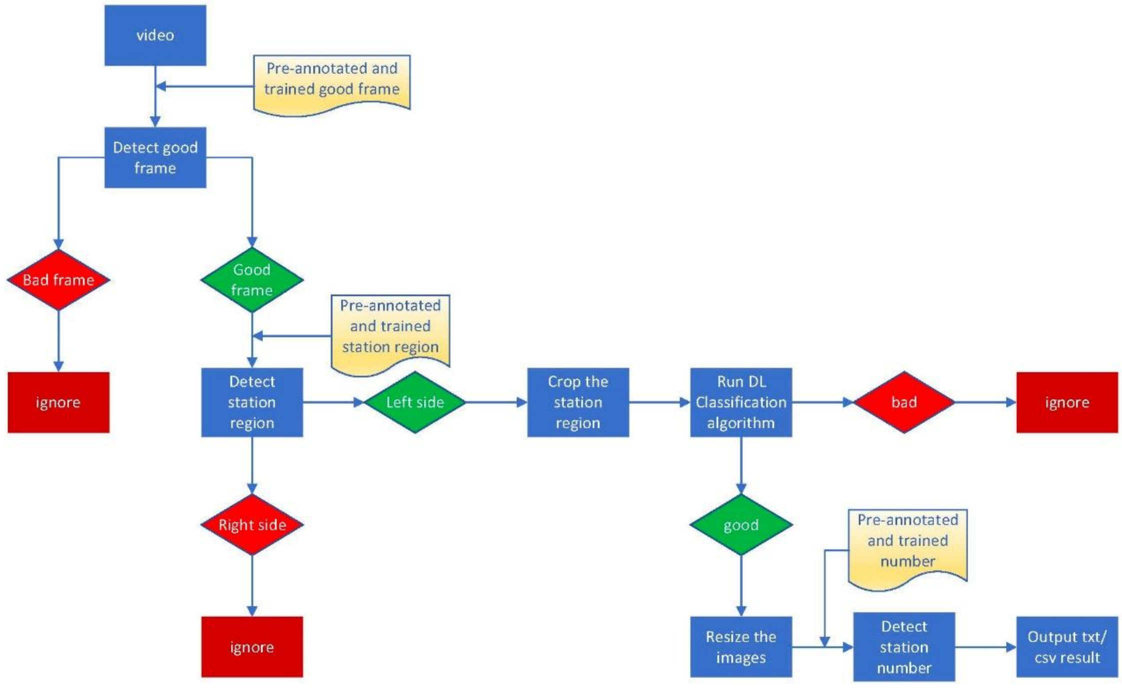

Figure 4. Flow chart of identification of a cow.

From raw position data (X, Y) to walking trajectories

This section describes the agent-based transmission model of Bovine Digital Dermatitis. This model is a modification of the original PeDViS-COVID-19 model (Atamer Balkan et al, 2024) to capture and model the key features of Bovine Digital Dermatitis. The PeDViS-COVID-19 model is a discrete time and space agent-based model that combines pedestrian movement with a microscopic COVID model to provide insight into indoor transmission dynamics. It has been successful in analysing the relative contribution of behaviors to transmission and the impact of control measures in reducing spread of COVID-19. PeDViS was designed to be reusable for other pedestrian systems (both real and modelled) and this is the first use case where the model has been changed from human movement to cattle movement in addition to a change in disease complex. An overall view of the model flow is shown in Figure 5.

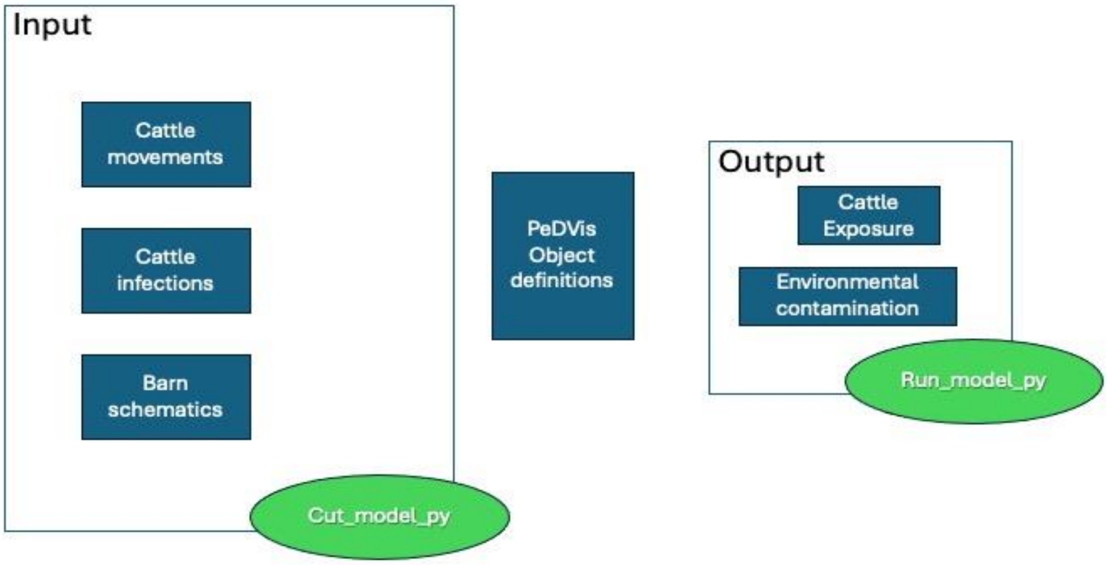

Figure 5. System Diagram, high level overview of the components of the data (blue boxes) transformation code and the model itself. The key code (green circles) is stored in cut_model.py and run_model.py, both use modified versions of the PeDViS system.

PeDViS-DD combines a cow movement model and an environmental spatially pathogen transmission model on which bacteria are shed into the environment and cows are exposed to bacterial contamination. Agent (cow) movement data has been derived from real-world tracking data from Dairy Campus using nine weeks of tracking data in the environmental barn for 193 unique individual dairy cows. There is a daily data gap of approximately one hour for each cow, corresponding to when they leave the barn for the milking parlour. The cow's infection status is determined weekly by a trimmer that scans the rear hooves of the cow for DD during milking. This data, which is granular to the day, is used to determine which cows are shedding contamination at any given time. For epidemiological analysis, these data can be used to determine the relationship between total pathogen contamination and future DD infection.

The PeDViS COVID-19 transmission model was converted to a version for bovine DD (PeDViS-BDD), taking the large set of cattle movement files (explained above, as it was the shared information from the tracking system [6] and transforming them into well-formatted agent (cows) scripts. In addition, a movement visualisation tool was created to visually validate the transformation of the movement files. All code is maintained in a GIT version control repository hosted by the WUR Data Competence Centre. The PeDViS-BDD outputs data in a format similar to PeDViS, that is a grid of contamination by coordinates over time and a list of contamination an agent (a cow) has been exposed to over time. The output data is stored as such to allow analysis of the results to be decoupled from the experimental runs. This decoupling is necessary due to the relatively long run time of the BDD model given the size of the cattle movement dataset.

One of the inputs to the model was a set of seventy different SQLite database files containing the daily movements of each cow in the research units of the environmental barn. The data are computed relative to a centroid in the barn itself and are granular to the second and to an accuracy of at least 10 centimetres. The movement component of the transmission model requires the definition of an *environment* in which the agents move and a set of *scripts* describing the movement of the agents. Both components were defined prior to runtime. The environment was constructed using an architectural schematic and an Excel spreadsheet defining the dimensions of the barn. The overall barn can be described as a series of barns into which groups of cattle are divided, and a central walkway connecting each stall to the main exit/entrance of the barn itself. There were no reports of changes in the environment between the start and end of data collection, so a single definition was sufficient for our experiments. The space of the barn is divided into grids of equal size that are used to represent the trajectories of the cows and the spatial unit of contamination when any giving grid is “walked” by an infected cow. If the grid was noncontaminated in the following time steps it will start the decaying process over time estimated by the decay function. An infected cow will spread bacteria over the grids that she passes during her trajectory at a rate determined by the shedding function. A susceptible cow that passes through contaminated grids will get exposure by the exposure function. The risk of infection is estimated in another model that will be explained below.

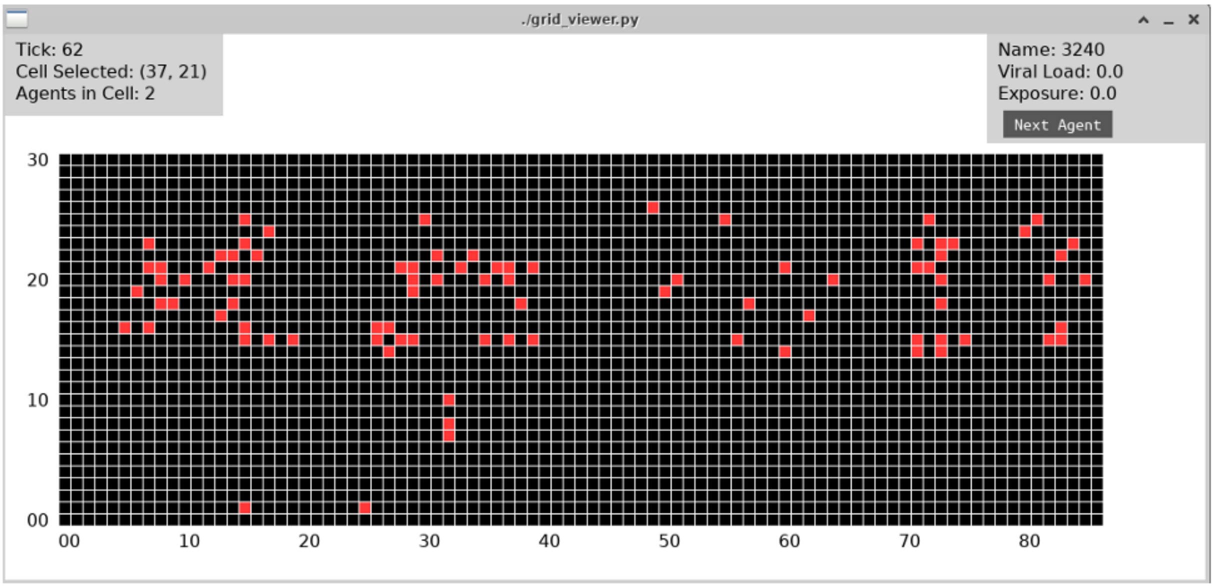

Figure 6. Grid Visualizer that displays an example of an agent location and movement in the mobility grid (small squares) at a given time in the barn space (black block).

Agent scripts are defined as a set of time-ordered actions per cow. In the movement dataset actions are limited to the following: enter, move, exit. Each action is accompanied by a set of Cartesian coordinates defining the location of the action. That is, except for exiting, which is assumed to have occurred at the previous set of coordinates for a given agent.

A Python script entitled cut_model was written to create both the environment definitions and the agents with their movement scripts. At a high level, the script first combines the movement files into a single database file for easier access and use. The script then determines the origin of the movement data; the captured data contains both negative height and width coordinates, while the PeDViS model requires all positive coordinates. All coordinates are adjusted off the new origin point to ensure all values are positive. After this, each cow's movement data is processed into a set of time ordered scripts. To reduce storage costs, if a cow does not move during a given time step (second) from their previous item in the script a new action is not recorded. Finally, each cow's infection status is calculated from a single file containing the infectious status of each hind hoof from the weekly scoring in the milking parlour by the trimmer. Given the size of this data set, extensive use of parallel computing was used in the cut_model script. It takes approximately one hour to generate seventy-four days' worth of minute granular movement across 193 unique agents (cows) on a 24 core Intel Xeon E5-2697 using 22 GB of memory. The raw movement data is stored in a single 27 GB sqlite database file. The produced agents and their scripts require 274 MB.

The basis of the PeDViS model is a set of Cartesian grids representing agent movement and environmental viral contamination. The pedestrian model directs agents through the mobility grid and transmission equations determine agent viral emission and uptake from the contamination layers. The grids themselves are rectangles defined by a width and a height in arbitrary units but calibrated to 10 cm movement cells. A mobility ratio is set before run-time to determine the relative granularity of agent space to contamination space.

As stated earlier, the general dynamics of bovine DD transmission are less understood than COVID-SARS-19, therefore the transmission component conversion required significantly more removal of features than the addition of features. The model only required the addition of a set of sub-location specific output files, the cow barns. The cows, when not in the milking parlour, are held in smaller groups in research units in this environmental barn. This typically means the cows are separated from one another and only overlap in space while entering and leaving the barn. However, occasionally, cattle returning from the milking parlour enter the wrong barn. Maintaining barn specific origin of cattle exposure allows for later infection attribution during analysis.

Epidemiologically, bacterial shedding, pickup, and decay onto the contamination layers each require their own expression as:

Shedding function: μ^2 / - 1+$e^{-\mu}$ + μ

Decay function: $e^{\{-\mu}e^{-\mu}$

Exposure function: $e^{\{\mu}e^{\mu}$

Further work will require the calibration of the parameters underlying these equations to the Dairy Campus farm environment. Every formula is expressed in terms of μ, simplifying the parameter search to a single dimension. Additionally, the granularity (time steps, grid cell sizes) of the model itself could be adjusted.

Finally, for quantifying transmission, we modelled environmental contamination using different equations to account each small spatial unit in the barn, using the observed spatial I(t) and standardising the shedding rate parameter to the spatial case. Infections are assumed to occur through the environmental compartment only. The environmental contamination used for each cow during the study was obtained as the output from the previous model.

We used a stochastic SIS stochastic compartmental model for estimating the transmission rate parameter (Chang and de Jong, 2023) and it is represented in Figure 7.

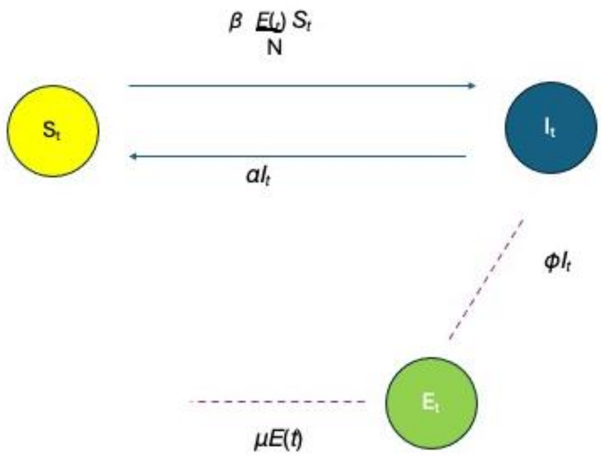

Figure 7. Representation of the stochastic SIS model with environmental transmission. St,It and E(t) represents the susceptible, infectious and environmental compartments, respectively. The solid lines represent the flow of individuals from one state to another. Transmission from St to It occurs at the rate β E(t)/N * St and the transition from It to St occurs at the rate αIt. The dotted lines represent the flow of pathogens in the environment, φIt represents the shedding rate at which infectious individuals shed the bacteria into the environment and μE(t) represents the decay rate at which bacteria decay in the environment.*

For transmission rate estimation, we used a binomial GLM model with cumulated E(t) for each recipient, cumulated over time and over the area covered by the cow in that time, as offset.

Locations of scripts

The scripts for monitoring locomotion with an examples of input and output for each step can be found in:

https://git.wur.nl/kbddht_pipelines/pipeline_locomotion

https://git.wur.nl/DairyCampus/NLAS/toolbox

Lessons learned

- A lot of time was spent developing the pipelines. However, processing Parquet files with pandas is quite slow, which will not perform well when taking the pipelines into production. We, therefore, recommend shifting most of the preprocessing tasks to the SQL Server. This should not require significant modifications to the pipelines since the development work has already been completed.
- During the course of this project, we moved from using a wearable tracking system like Tracklab as primary technology to assess behaviour and using cameras to validate this tracking system to proposing a multi-camera system as major tool to assess behaviour through tracking of 3D key points assigned to bounding boxes, and the tracking system as a potential bounding box identifier. There were multiple reasons for this switch (some mentioned below), and included the realisation that more detailed and direct information can be obtained with computer vision;
- The Sewio system is currently not fit for sampling at high resolution for a long period of time without recharging the battery. When collecting data at high resolution (20Hz), the battery lifetime reduces significantly to two weeks. This would require a lot of manual work with the risk of making errors, and a high risk of losing data;
- 3D tracking is not possible with the current tracking system;
- The Sewio system does not automatically monitor whether raw data is generated and stored, nor does it automatically warn researchers in case the data storage capacity exceeds its limits. This increases the risk of losing raw data;
- To monitor the locomotion of cows, video footage of cows walking past the camera without interruption is required. When creating a camera-setup ensure that disturbances are minimal;
- Inaccuracies in cow identification in the exit lane of the milking parlour were assigned to the substantial physical distance between the point where cow identification is read by the selection gates and the point where the cow is passing the camera view. With that, cow identification is still work in progress. Practical usefulness of continuous monitoring is dependent of this cow identification;
- Camera calibration and a multiple camera set-up have advantages for processing raw footage.
- Pre-installation preparation of hardware can be improved considerably. We recommend to calibrate camera's before installation, preferably in a setting that is optimized for this task. For further details on calibrating camera's and a multiple-camera set-up please see the following document;
- Security cameras (as usually applied for vision applications in livestock barns) in general have a width of the field of view that is too large for standard camera model calibration. Despite that theoretically fisheye cameras do not have one unique point where all image lines cross (in contrast to pinhole cameras) assuming such a point for security cameras is feasible and enables the calculation of stereo vison.
- Camera positions and orientation of cameras that are used to obtain stereo vision should be checked (and updated if needed) regularly, particularly after lens cleaning.
- After annotation of images for model training check the results for potential inconsistencies.
- Annotation of video footage for model development can best be done iteratively. This means, first start with annotating a limited set of selected video footage capturing as much variation as

possible. In our case for example different cow coat patterns, different lighting conditions and different viewpoints. Then assess the model performance, and find out where accuracy is lacking. Annotate more data if needed to improve accuracy and repeat this iteration if necessary.

- Long term collection of raw position data with high temporal resolution with standard tracking sensors is considerably more difficult than expected. Battery capacity was limited (less than 3 weeks when XY-positions were collected at 10 Hz). As a consequence human intervention was required often to replace empty tags with full tags. Recording of combinations of tag identification and cow identification is essential for the interpretation of the data, but is a none automated procedure. Due to the required human intervention, this is error prone which limits useability under practical circumstances (on commercial farms).
- Accelerometer data collected with sampling rates of 1 Hz or less probably has limited useability to obtain information about behaviour because then specific movement patterns characterising behaviours (e.g. eating, ruminating, walking) presumably are not well captured.
- Computer vision applications to detect specific lesions and disease stages of the claws under commercial operations presented several challenges and need more time and efforts to be optimized;
- Hardware adequacy needed for working conditions on commercial farms should be optimized and long-term use tested;
- Maintenance and lens cleanliness protocols for the cameras, particularly in the milking parlour, should be optimized and adapted to many different farm conditions and settings. Cameras located here have to function in an extremely harsh environment (moist, water, manure, urine, dust). In addition, if the camara is expected to be used for longer periods (1 month or more) it should be considered to protect and seal them at the level of an underwater camara, to avoid corrosion of cables and connections and condensation, that either affects the quality of the image or the functioning of the camaras;
- The developed camera system set-up for automatically detecting DD was satisfactory for a rotary milking system, but the approach used, and protocol developed should be adapted for other milking systems (such as herringbone or robot systems) on commercial farms;
- Although quality of the images from surveillance camaras was good enough, the use of industrial camaras should be explore since these can be more flexible in their settings and could have more clarity;
- For DD, a more extended data set is required, specially of sick animals to increase the spectrum of clinical presentation, signs and variability.

List of people involved and their role

Name	**Email address**	**Role in project**
Claudia Kamphuis	Claudia.Kamphuis@wur.nl	Senior researcher at Animal Breeding and Genomics. Researcher. Project lead.
Wijbrand Ouweltjes	Wijbrand.Ouweltjes@wur.nl	Senior researcher at Animal Health & Welfare. Researcher. Data analysis. Welfare expertise.
Bert Klandermans	Bert.Klandermans@wur.nl	Data engineer at Dairy Campus. Raw data collection and data quality control.
Tomas Izquierdo		Researcher at Animal Health & Welfare. Modelling. Computer vision expertise.
Marjaneh Taghavi	Marjaneh.Taghavi@wur.nl	Researcher at Animal Breeding and Genomics. Modelling. Computer vision expertise.

Ines Adriaens	Ines.Adrianes@wur.nl	Researcher at Animal Breeding and Genomics. Modelling. Data scientist.
Joop van der werf	Joop.vandenWerf@wur.nl	Research technician at Animal Health & Welfare.
Gustavo Monti	Gustavo.Monti@wur.nl	Associate Professor at Quantitative Veterinary Epidemiology. Project lead. Epidemiology of DD.
Mart de Jong	mart.dejong@wur.nl	Professor at Quantitative Veterinary Epidemiology. Project lead and modelling.
Piter Bijma	Piter.Bijma@wur.nl	Associate Professor at Animal Breeding and Genomics. Project lead. Genetics.
Zhen Pan		Researcher at Quantitative Veterinary Epidemiology. Computer vision expertise.
Colin Teberg		Steady State Scientific Computing, Chicago, Illinois, United States of America- Software Computing.

Output

Scientific output:

Taghavi M., H. Rusello, W. Ouweltjes, C. Kamphuis, I. Adriaens. 2023. Cow key point detection in indoor housing conditions with a deep learning model. Journal of Dairy Science. DOI: 10.3168/jds.2023-23680

Presentations:

Kamphuis C., 2023. Continuous and non-invasive monitoring of animals using sensors and computer vision analysis. Wageningen Data and Modelling day, 16 November 2023, Ede, the Netherlands.

Kamphuis C., 2023. Data een blik in de toekomst. Data Dairy Dag bij Dairy Campus, Leeuwarden, September 19, 2023

Koning de, C., C. Kamphuis, W. Ouweltjes. 2022. Digital Dairy Developments. 2022. World Agri-Food Innovation, Beijing

Taghavi M., Russello H., Ouweltjes W., Kamphuis C., Adriaens I. 2022. Automated gait analysis with a deep learning key point detection model. 10th European Conference on Precision Livestock Farming (ECPLF), Vienna, Austria.

Taghavi, M., C. Kamphuis, T. Izquierdo, I. Adriaens. 2023. Automated gait analyses with a deep learning key point detection model. 5th Precision Livestock Farming seminar, Copenhagen, Denmark.

Ouweltjes, W. 2022. Next level animal science: monitoring behaviour. Webinar Japan-Netherlands on digital dairy farming & Animal welfare assessment.

Gustavo Monti, Piter Bijma, Quirine ten Bosch, Mart C.M. de Jong. Modelling Environmental Transmission of Digital dermatitis in Dairy Cattle using sensor and computer-vision derived data. Accepted, to be presented during 17th International Symposium on Veterinary Epidemiology and Economics, Sidney, Australia, November 2024.

Popular press:

Continue monitoren op gezondheid loont. Janet Beekman. 2023. Boerderij 108 – no. 50 (13 september 2023)

Cameras for healthy hooves. 2023. NLAS Magazine. https://magazines.wur.nl/nlas-magazine-en/cameras-for-healthy-hooves

Smart Technology helps track and improve cow health. 2022. https://www.youtube.com/watch?v=V_FPqSx0Upk

De technologie houdt een oogje in het zeil. Een zesde zintuig voor de boer. Nienke Beintema. 2022. WageningenWorld|3. https://edepot.wur.nl/582933

Camera's verraden de gezondheid van de koe. Nienke Beintema. 2022. Resource. https://edepot.wur.nl/582175

Onderzoekers Dairy Campus willen kreupele koeien via camera's vroeg herkennen. Guus Daamen.2022 Melkvee. https://www.melkvee.nl/artikel/459072-onderzoekers-dairy-campus-willen-kreupele-koeien-via-cameras-vroeg-herkennen/

Een camerasysteem detecteert kreupele koeien automatisch. 2023. Een camerasysteem detecteert kreupele koeien automatisch – Dairycampus

Kamphuis C. 2023. Non-invasief langdurige monitoring van gezondheid en welzijn van een groep dieren – gedrag gerelateerd aan locomotie en transmissie van Digitale Dermatitis met tracking data en beeldverwerking. DairyCampus Innovatiefonds factsheet