High Classification Rates for Continuous Cow Activity Recognition Using Low-Cost GPS Positioning Sensors and Standard Machine Learning Techniques

Torben Godsk[1,2] and Mikkel Baun Kjærgaard[2]

[1] Knowledge Centre for Agriculture, Agro Food Park 15,
DK-8200 Aarhus N, Denmark
tbg@cs.au.dk

[2] Department of Computer Science, Aarhus University, IT-parken, Aabogade 34,
DK-8200 Aarhus N, Denmark
{tbg,mikkelbk}@cs.au.dk

Abstract. In precision livestock farming, spotting cows in need of extra attention due to health or welfare issues are essential, since the time a farmer can devote to each animal is decreasing due to growing herd sizes and increasing efficiency demands. Often, the symptoms of health and welfare state changes, affects the behavior of the individual animal, e.g., changes in time spend on activities like standing, lying, eating or walking. Low-cost and infrastructure-less GPS positioning sensors attached to the animals' collars give the opportunity to monitor the movements of cows and recognize cow activities. By preprocessing the raw cow position data, we obtain high classification rates using standard machine learning techniques to recognize cow activities. Our objectives were to (*i*) determine to what degree it is possible to robustly recognize cow activities from GPS positioning data, using low-cost GPS receivers; and (*ii*) determine which types of activities can be classified, and what robustness to expect within the different classes. To provide data for this study low-cost GPS receivers were mounted on 14 dairy cows on grass for a day while they were observed from a distance and their activities manually logged to serve as ground truth. For our dataset we managed to obtain an average classification success rate of 86.2% of the four activities: *eating/seeking* (90.0%), *walking* (100%), *lying* (76.5%), and *standing* (75.8%) by optimizing both the preprocessing of the raw GPS data and the succeeding feature extraction.

1 Introduction

Due to intense competition in the domain of precision livestock farming, the farmers need assistance from either qualified extra man power or modern technology to overview and attend the herd, in order to effectively find *focus cows* that for some reason needs special attention or relief care. It requires full attention to do so, in order to prevent false positives or, what may be even worse, overlooking a true positive causing an animal to suffer.

P. Perner (Ed.): ICDM 2011, LNAI 6870, pp. 174–188, 2011.

A global navigation satellite system, like for instance the Global Positioning System (GPS), is a widely used technology for various position based applications. The main reason for considering GPS for monitoring cows is that locally the positioning technology is infrastructure-less, in contrast to an alternative like a local sensor network, e.g., as used by Nadimi et al. [9]. However, an infrastructure may be required for comunication.

Pattern recognition and machine learning are widely used techniques to recognize patterns in data. In this paper we present high classification rates obtained by preprocessing position data and extracting a broad variety of features that serve as input to standard machine learning algorithms for classification of specific cow activities. The classification results outputted by the standard machine learning algorithm are optimized by adjusting the input features, i.e., adjusting the preprocessing of data as well as the succeeding feature extraction.

The dataset used to evaluate the proposed method combines continuous position data from 14 dairy cows on grass rigged with low-cost GPS receivers with continuous manual observations of the cows' activities. We use position data from more than just one cow since the individual cows have a tendency of behaving differently and finding their own routines in their way of performing their activities as described by Phillips et al. [10]. We therefore subdivide cow behaviors into activities with cross cow commonalities and classify the behavior with regards to these. A restriction in our study is, that we only consider the activities independently and not the transitions between them. The individual combination of activities defines each animal's normal behavior. The goal is that the activity recognition can be used to observe when an animal start behaving abnormally, i.e., when the activities performed diverges from the normal behavior of the individual cow, since it often indicates a change in the state of health and/or welfare.

2 Related Work

Previous research shows that feed intake depends on a cow's health condition [3], and time spent at the feeding area correlates with feed intake [4], moreover, abnormal lying behavior correlates with lameness amongst cows [5].

A study, by Agouridis et al. [1] examines GPS collar capabilities and limitations in regards to tracking animal movement in grazed watersheds, conclude that the position accuracy decreases as cows move under a tree or so, and thereby loose *line of sight* towards the GPS satellites. That GPS performance degrades in terms of both coverage and accuracy when experiencing problematic signal conditions due to attenuation is analyzed by Kjærgaard et al. [6].

Schwager et al. [12] measure cows' moving speed via hi-end GPS receivers. In addition they measure head roll and head tilt with accelerometers. They apply the measurements to a simple K-means classification algorithm without *a priory* information. This leads to a repeatable categorization of the animals' behaviors into periods of activity and inactivity. Though, using hi-end GPS receivers would give better position quality, we use low-cost GPS receivers in an attempt to meet the basic requirements of scalability when monitoring a bigger herd.

Nadimi et al. [9] use a local ZigBee based sensor network to track and classify cow behavior. They too derive the moving speed, head roll and head tilt, and by using a simple classification tree they too succeed to classify both activity and inactivity. In comparison, our approach has limited maintenance due the infrastructural independence. In addition, better scalability is achieved firstly, as there are no upper limit for neither the number of receivers nor the size of the area being monitored and secondly, achieving good results by using low-cost receivers in the experiments instead of hi-end equipment, makes monitoring of bigger herds affordable.

Robert et al. [11] use three dimensional accelerometers and video based observations for classifying behavior patterns in cattle, and classify lying (99.2%), standing (98.0%), and walking (67.8%). In comparison, we manage to recognize the activity of a cow walking in 100% of the occasions. However, we are unable to match the succes rates of both lying (76.5%) and standing (75.8%), which indicate that introducing other types of sensors, e.g., an accelerometer, might improve our results. In addition and unlike their work, we recognize the activity of a cow eating and seeking and obtain a succes rate of 90.0%.

3 Collection of Position Data for Cow Activities

The GPS receivers used for the experiment are *i-gotU GT-600* − a commercial low-cost receiver [7] with a *SiRF Star III Low Power* chipset scheduled to log a GPS position every second. The receivers were installed in a plastic housing as depicted in Fig. 1(a), and mounted on the cow collars as illustrated in Fig. 1(b).

The 14 cows used in the experiment are arbitrarily picked out from a herd of 28 dairy cows, i.e., they were selected with no regards to their expected behavior during the experiment. The reason for using 14 cows instead of the entire herd is based in the practical challenge in observing the animals manually, while taking useful and trustworthy notes to be used as ground truth in the analysis. Though, the observers where stationed at static observation points using field glasses to watch the animals from a distance, it was unavoidable to disturb the animals in some sense, as the observation points had to be in the middle of each of the two consecutive fields in order to guarantee visual contact with the animals at

(a) *Receiver and housing*

(b) *Mounted receiver*

(c) *Manual observations*

Fig. 1. Setting up an performing the experiment

all times, as Fig. 1(c) depicts. However, our focus in this work is on recognizing particular cow activities in contrast to recognizing the normal behavior of each cow. Therefore, should a cow stop and stare at an observer for a while, the sequence is simply annotated as the cow performing the standing activity.

Organizing Data. The different behavioral changes, that may be used to point out a possible *focus cow* candidate are many. Typically, the normal behavior of cows in a herd diverges from one animal to another. However, by detailing the behavior into lower levels of cow activities like: *walking*, *lying*, *standing* and *eating seeking*, the cows' way of performing these detailed activities becomes similar. From the duration, combination, frequency, etc. of these detailed activities performed by any normal behaving cow, it would be possible to define its normal behavior. In this work, we strive to recognize such immediate cow activities, to assist a domain expert in the work of detecting behavioral abnormalities amongst cows. In order to meet the cross-animal physiological variations, the many different behaviors are divided into common activities of a lower level of abstraction, which all may serve as abnormal behavior indicators, e.g. jumping, toddling, lying and eating. However, using a low-level GPS receiver cause some limitation in terms of the position information provided. The information is limited to: time of measurement, latitude, longitude, elevation, and speed. In addition, the sample rate has a maximum of 1 sample per second. Therefore, not all indicators are detectable from the provided position data, i.e., the activity has to affect the movement taking place from one measurement to another. This excludes indicators like toddling and jumping, and leaves a subset of activities detectable when using position data. From this subset we define four activities to look for:

Walking defines the activity of the cow walking towards a goal, e.g., from A to B without stopping or simply tagging along other cows. Should the cow stop for any reason, it is no longer considered to be walking. This often takes place when the cow moves from one field to the consecutive one, or when the cow moves to the drinking vessel.

Eating seeking defines when a cow shows eating behavior, i.e., it either eats or seeks for grass, and possibly the cow stops from time to time chewing. The *eating/seeking* activity is the hardest to recognize, since the cow either is *walking* around seeking for grass or *standing* still eating, and therefore tends to be confused with the other activities.

Standing defines when a cow stands still for a longer period of time, e.g., thirty seconds or more without showing neither *eating* nor *seeking* behavior. It may be hard to distinct this activity from *lying*. However, a standing cow tends to be moving just a little more than one lying down, causing the measured position to move in contrast to a cow lying still.

Lying defines when a cow lies down for a longer period of time, e.g. thirty seconds or more. When the cow lifts its head and looks around this activity may easily be confused with *standing*.

Selecting data sequences. Sequences of data where cows are doing one of the four activities were handpicked from the full dataset. Any sequence selection is based upon a manual observation of high quality stating that the cow is performing an activity of interest. The sequences are selected in a manner so that they together represent all 14 cows doing all 4 activities of interest for a period of at least 4 minutes. In that way we get a dataset where each cow performs each activity at least once and for a minimum of four minutes.

Due to the data sequences selection strategy there may be both time and distance gaps between two sequences. To prevent these gaps from influencing on the results, the sequences are treated as atomic datasets instead of one assembled dataset with adjacent sequences, i.e., the last measurement from the previous sequence is discarded when loading a new sequence. A drawback of this approach is that the transition between two activities is neglected. In this work it is considered a trade off in order to work with noise free data, however, we will consider this issue in our future work.

4 Recognition of Cow Activities

We approach the activity recognition problem from a software perspective and leave the classification to a machine learning toolkit in this case the Weka Toolkit [14]. We present a method for obtaining the highest classification success rate by optimizing the preprocessing of raw GPS position data and the extraction of features that serves as input to the machine learning algorithm, instead of optimizing the machine learning algorithms themselves.

Figure 2 shows how the activity recognition module is divided into three analyzing blocks: (1) the Movement Analyzer (MA) process the raw GPS position data, determines the movement taking place between two adjacent measurements and represents it in a Movement Data Structure (MDS); (2) the Segment Analyzer (SA) groups the MDSs into segments of a certain size, all the movement information are processed and as a result a broad variety of features are extracted and represented in a Segment Data Structure (SDS); (3) the Activity Analyzer (AA) use the SDSs as input to the machine learning algorithm and represents the classified activity in an Activity Data Structure (ADS).

Designing the module with three analyzers each with different data structures as output is to some extent inspired by research done by Zheng et al. [15] where they recognize commuters' different transportation modes like walking, bicycling and driving from raw GPS data. They assemble a number of measurements in segments, which are constituted by a starting point where the current mode of transportation is initiated and an ending point where the transportation mode changes. By extracting numerous features e.g. *heading change rate* from the GPS data within such a segment and processing these features via machine learning, they extract the information from raw GPS data through data mining without using neither *a priori* information nor on-time user inputs - except from information on where the transportation mode changes takes place.

Each of the three analyzers consume and produce specific data structures; the data structures are illustrated as white squares on the right side of Fig. 2. Each analyzer is individually adjustable so that the optimal feature extraction

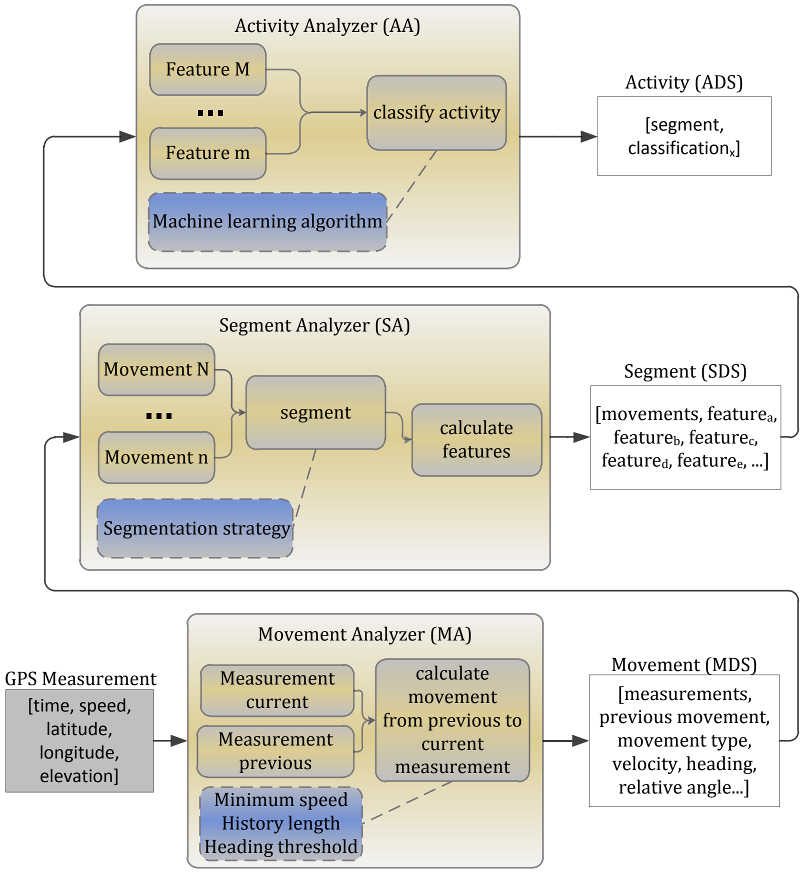

Fig. 2. The activity recognition module with three analyzers, the data structures to the right, and blue boxes illustrating the input parameters for adjusting the feature extraction.

can be obtained causing the best activity recognition results; the dashed boxes show the specific parameters used for adjusting the individual analyzer. The grayed square on the bottom-left side of the figure shows the incoming GPS measurement provided by the low-cost GPS receivers. The three analyzers and their corresponding data structures are described in details in the following.

Analyzing movements. The main goal for the MA is to extract information about, what happen between two adjacent GPS measurements, e.g., how far did the cow travel from measurement$_{t-1}$ to measurement$_t$, at what velocity and in which direction. The MA takes incoming GPS measurements and produce MDSs, which represent the relation between the last two received measurements. The *movement* makes the foundation of the succeeding feature extraction.

Besides providing information like: speed, acceleration, absolute heading, and distance traveled, a discrete representation of the movement is categorized as being either: left turn, right turn, forward, u turn, or non moving. This approach is inspired by the domain of various sports, where athletes' movement patterns have been used to analyze their physical fitness and performance. Mohr et al. [8] used such a discrete representation when classifying activities into standing, walking, jogging, sprinting etc. to analyze the performance of high-standard soccer players. Also Spencer et al. [13] analyze elite field hockey players' performance during a game by filming the players during the game and discretize the time-motion information into movements for classification into more or less the same discrete representations as above. A condensation of the information contained by a MDS is listed in Table 1.

Table 1. A condensed list of information regarding each individual movement

Parameter	Description	Example
movement type	a discrete representation of the latest type of movement performed	left, right, forward, u turn, non moving
angle	angle relative to the previous movement	[deg] and [rad]
magnitude	distance traveled between the two measurements	[m]
speed	estimated speed	[m/s]
heading	absolute heading	[deg]
acceleration	based on estimated speed of the last two measurements	[m/s^2]

As illustrated in Fig. 2, the MA can be adjusted via one or more of three input parameters. Whether a cow is moving or not is determined using a naive Bayesian filter, as illustrated with pseudo code in Fig. 3. Each of the three input parameters have different influence on the MA: *minimum speed* defines the threshold between non moving and moving; *heading threshold* defines the threshold for whether the current movement type is forward, left, right or u-turn; and finally *History length* is the number of old movements taken into account in the Bayesian filter when deciding whether the current movement type is moving or non moving. The selection of the likelihood$_{nonmoving}$ constants 0.1 and 0.6 in the pseudo code is based on experience from previous lab work with detection of bicyclist and pedestrian movements.

Selection of the MA input parameter values was based on intuition and experiences from observing cow behavior. Hence, the *minimum speed* was set to 0.3 m/s as cows walking towards a certain goal tends to move at that pace or faster, *history length* was set to 4 by pure intuition and *heading threshold* was set to 40 degrees for the same reason.

$prior_{nonmoving} = 0.5$
$prior_{moving} = 1.0 - prior_{nonmoving}$
$accuracy_{GPS} = 0.4$

$$uncertainty = \begin{cases} 1.0 & \text{if } distance_{movement_t} < accuracy_{GPS} \\ \left(\frac{accuracy_{GPS}}{distance}\right)^2 & \text{else} \end{cases}$$

for each $movement$ in $list_{history}$:
{

$$likelihood_{nonmoving} = \begin{cases} 0.10 & \text{if } minSpeed \leq v_{movement} \\ (0.60 \times uncertainty)^2 & \text{else if } uncertainty < 1.0 \\ 0.60 & \text{else} \end{cases}$$

$$likelihood_{moving} = 1.0 - likelihood_{nonmoving}$$

$$p_{nonmoving} = \frac{prior_{nonmoving} \times likelihood_{nonmoving}}{\left(prior_{nonmoving} \times likelihood_{nonmoving}\right) + \left(prior_{moving} \times likelihood_{moving}\right)}$$

$prior_{nonmoving} = p_{nonmoving}$
$prior_{moving} = 1.0 - prior_{nonmoving}$
}
$isMoving = prior_{nonmoving} < 0.5$

Fig. 3. Pseudo code for determining motion for a movement instance

Analyzing segments. The main purpose of the SA is to extract features from movements and pass the feature information on as segments. The SA assembles incoming movements, and extracts a broad variety of features from the movement assembly, once a certain number of movements has been assembled. The criteria for segment completion is customizable via the *segmentation strategy* parameter, as depicted in Fig. 2. Depending on the domain usage, such a segmentation strategy may vary, e.g., segment when the timespan between the timestamps of the first and the last measurement reaches a certain limit.

With inspiration from research done by Zheng et al. [15], we extract a broad variety of features from the *movement* data, e.g. *HeadingChangesDegreesForwardRate*, which represents the rate of heading changes in degrees for all movements in the segment moving straight forward.

The SA computes fifty six different features represented in the SDS. As many of these features tend to be variants of each other, they are grouped for clarity and listed in Table 2.

Selecting the segmentation strategy to use as input parameter for the SA was based on experiences from observing cow behavior. We found that cows often do the same activity for two to three minutes or more, hence, *segmentation strategy* was set to segment every 160 seconds. Also we assume that too small a timespan might lead to large variations between the segments. However, the segmentation strategy remains to be tested properly, as we in this work omit to consider the transition between activities.

Table 2. A list of features extracted for each segment

Parameter	Feature
movement type	Distribution (% of forward, left, right, etc.)
	Change rate between moving and non moving
	Change rate between any type of movement
heading	Changes accumulated (forward, left, etc.)
	Change rate (forward, left, etc.)
	Changes max (forward, left, etc.)
speed	Max, min and mean
acceleration	Max and min
	Mean and accumulated (both positive and negative)
	Changes between positive and negative
distance	Accumulated for 2D and 3D (moving and non moving)
	Max for 2D and 3D (moving and non moving)
time	Accumulated (moving and non moving)

Analyzing activities. The AA is responsible for processing the incoming features provided via the SDSs, and classify the current activity using a machine learning algorithm. The segments are processed by providing the incoming feature instances to standard machine learning techniques implemented by the Weka Toolkit [14]. We approach the cow activity recognition problem from a software perspective, hence we use a standard machine learning API, and omit to optimize the machine learning algorithms and techniques. The *machine learning algorithm* parameter provides the ability to change the algorithm used, as illustrated in Fig. 2. The classification result and the corresponding SDS are represented in an ADS. Based on experience from previous lab work with classification of bicyclist activities, we used a random classifiers committee (END) as *machine learning algorithm* in this work. In addition, we compare these results with the classification rates of other well performing algorithms.

5 Results

The relevant sequences of data selected for the following analysis consist of position tracks where a cow performs one of the following activities: *lying*, *walking*, *eating seeking* or *standing*, as described in Sect. 3. The sequences sums up to a total of 16 hours of unbalanced data, where the *walking* activity represented by 50 minutes of data is the one activity with the least data available, followed by *lying* with 136 minutes, *standing* with 165 minutes and finally *eating seeking* with 613 minutes.

Setting the configuration parameters initially was based on intuition and experiences from observing cow behavior. Consequently, the input parameters where set as follows: *minimum speed* was set to 0.3 m/s, *history length* was set to 4, *heading threshold* was set to 40 degrees, *segmentation strategy* was set to segment every 160 seconds and finally an END random classifier committee was selected as *machine learning algorithm*. The background for selecting these features is explained in Sect. 4. By using this parameter configuration we obtain a success rate of 86.2%, and the result was evaluated against various combinations of different input parameter values as listed in Tables 3 and 4.

Table 3. The MA and SA input parameter values used for the evaluation

Parameter	min	max	step
movement analyzer (MA):			
minimum speed [m/s]	0.1	0.5	0.1
heading threshold [deg]	10	50	10
history Length	1	10	1
segment analyzer (SA):			
segmentation strategy:			
- Timespan in seconds	30	180	10

The individual classification rates of one thousand iterations of each combination evaluated using an END random classifier committee reaches from 71.8% to 86.5%. The END classifier is used with its default configuration, i.e., 10 committee members and is evaluated using 10 folds cross validation. In addition, we tested several machine learning algorithms also provided by the Weka Toolkit [14], and in Table 4 we present the results of the best performing ones having set the input parameters as stated above. We found END to be best performing in terms of average success rates. In addition, the table shows the mean time of processing one instance after running the one thousand iterations of 361 instances on a Intel(R) Core(TM)2 Duo CPU T8300 (2.40GHz,2.40GHz) with 3.00 GB RAM, on a 32 bit Windows 7 Enterprise operating system.

Table 4. Results of the evaluation of the algorithms performing best in the test

Machine learning algorithm	avg success rate %	milliseconds/instance
- END	86.2	5.5
- SMO (SVM)	85.7	10.3
- Classification Via Regression	85.7	13.8
- Random Forest	85.5	1.8
- J48	85.4	1.0

By evaluating the results of the END based classifications, with only one of the four input parameters varying at a time, we found, that selection of any of the tested values for *minimum speed*, *heading threshold* and *history length* has very little impact on the success rate for the given data.

In contrast, the results of testing the segmentation strategy shows a raising tendency of the classification rates as the segment size increases, as depicted in Fig. 4. However, the graphs seems to stagnate after reaching the selected strategy, where the timespan between the timestamps of the first and the last measurement is 160 seconds. The same characteristics tends to match all the algorithms listed in Table 4.

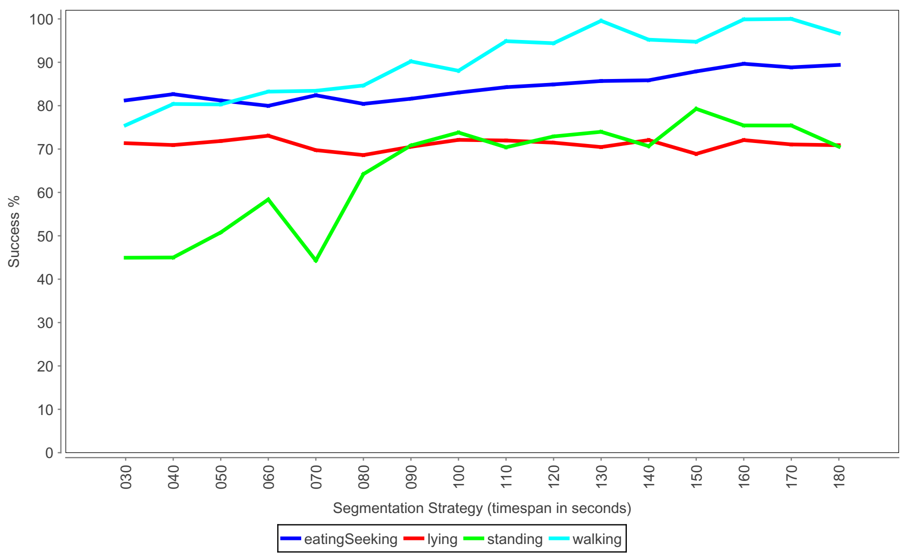

Fig. 4. Graph showing the configuration with variable segmentation strategy

For some features, an inspection of their characteristics leads to an explanation of the mutual difference in success between the four activities. For instance, Figure 5 depicts the distribution of the heading change rate in degrees while the cow is moving forward. The figure indicates that the *walking* activity dissociates itself from the other activities, which may explain the success rate of 100%. Moreover, the figure shows the same tendency for approximately 2/3 of the cases of performing the *eating seeking* activity. However, for the 30% fractile it tends to hide behind the *lying* distribution graph. In addition and similar for all feature distributions, it looks like this feature is of no help in the distinction between *lying* and *standing* activities as they collide for almost all values of the heading change rate in degrees while moving forward. It explains by the fact that the three activities are composed by either none or limited forward movements in contrast to the *walking* activity. The search for features making the distributions diverge will be challenged in future work.

The number of false classifications exposed in the confusion matrix in Table 5 verifies, that both the definitions of the activities of *lying* and *standing* are similar, and that the definition of the *eating seeking* activity cause for it to be confused with the for two activities, due to the many and long periods of time where the cow is standing still chewing and grassing.

Table 6 sums the number both false positives and false negatives and lists the success rates. In the domain of precision livestock farming the number of both false positives and false negatives are severe, as they may lead to animals to suffer unattended. Moreover, the combination of false positives and negatives is important, as an eating and seeking cow classified as a lying cow, will appear as if the feed intake is decreased and the resting activity is increased, which may indicate the cow as being in need of extra attention, i.e., a *focus cow*.

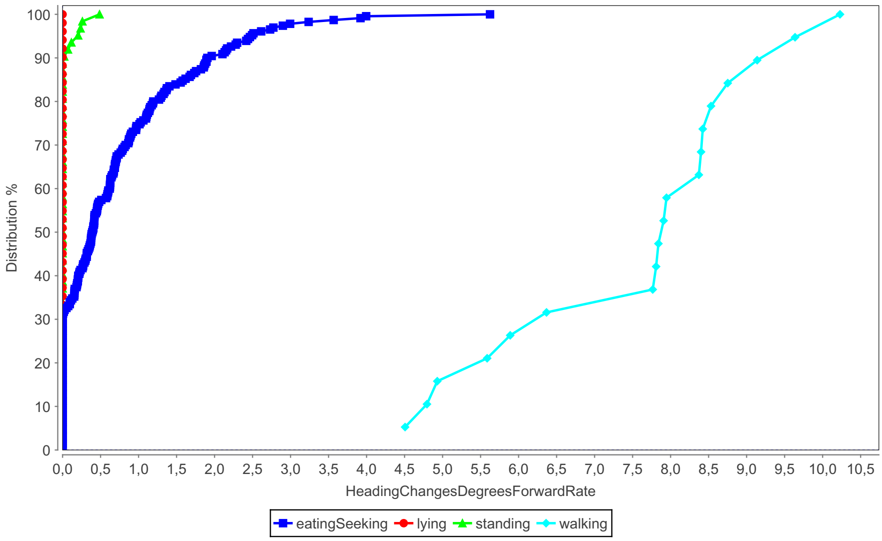

Fig. 5. Distribution of *heading change rate in degrees while moving forward*

Table 5. Confusion matrix of the cross validation

classified as:	lying	standing	walking	eating seeking
observation:				
lying	76.5% (39/51)	13.7% (7/51)	0.0%	9.8% (5/51)
standing	3.2% (2/62)	75.8% (47/62)	0.0%	21.0% (13/62)
walking	0.0%	0.0%	100% (19/19)	0.0%
eating seeking	3.0% (7/230)	6.5% (15/230)	0.4% (1/230)	90.0% (207/230)

Table 6. Summed false negatives and false positives

	false negatives	false positives	success
lying	23.5% (12/51)	2.9% (9/311)	76.5% (39/51)
standing	24.2% (15/62)	7.3% (22/300)	75.8% (47/62)
walking	0.0%	0.3% (1/343)	100% (19/19)
eating seeking	10.0% (23/230)	13.6% (18/132)	90.0% (207/230)

Summary. We find that varying the input parameter configuration has very little impact on the given dataset. For instance, the tested values and combinations of both *history length*, *minimum speed* and *heading threshold* has very little influence on the classification rate. However, the *segmentation strategy* used in this work shows a tendency of an increasing success rate as the segment size increases. Moreover, it is not thoroughly tested as the detection of transitions between activities are omitted in this work.

6 Discussion

In this section we will discuss improvements in hardware, collection of additional datasets and the recognition of abnormal behavior.

Improving the Hardware. An issue with the low-cost GPS receiver used for this experiment is that they stop calculating new GPS positions to save power after remaining still for an unspecified period of time. This functionality is unmanageable when using the GPS receiver's scheduling mechanism, however it seems to become an advantage for the machine learning algorithm in the distinction between the two activities of *standing* and *lying*, as the GPS positions seems to drift a little when a cow is standing still in contrast to when a cow is lying still. However, for these two activities we are unable to match the results by Robert et al. [11], which indicates that we might improve our success rates by fusing measurements from other types of sensors like an accelerometer with the low-cost GPS receiver. Given the achieved classification rates using a low-cost GPS receiver as sensor, one can only expect even better classification rates in the future as the existing positioning technologies mature and new promising global navigation satellite systems like Galileo [2] becomes operational.

Collection of Additional Datasets. This work was based on sequences of data with manual observations as ground truth. The benefit of manual observations is that we were able to monitor fourteen cows moving around freely over two consecutive fields. Given fourteen animals also means that physiological variation is represented in the data which will decrease along with the number of animals. However, the manual method also limits our dataset because we had to select particular tracks from it which may cause the activity recognizing model to be trained and tested with less noisy data. Moreover, we treated each of the selected sequences of position data as atomic datasets. An obvious approach for future work would be to use datasets including transitions between activities, e.g. a dataset where a cow after *standing* still for a period of time, it *walks* until it reaches a location, where it starts to *eat* and *seek* for a while before *lying down*. This also enables us to apply, e.g., a hidden Markov model to model the transitions among the activities over time.

Another method for capturing ground truth would be to use video recording to document ground truth, especially, this would remove uncertainty in situations where the manual observation diverges from the position data, e.g., if a cow is observed to be lying down while the position data reveals that the cow is actually moving around. However, video recording both limits the number of animals that can be monitored and the size of the field to keep the animals in view. Therefore given the same human effort the video-based method can produce data for fewer animals thereby decreasing the physiological variation. In our future work we plan to experiment with introducing video based observations because this would enable us to better study transitions between activities which is difficult to capture accurately with manual observations.

Recognizing Abnormal Behavior. Recognizing specific cow activities is the first step towards spotting *focus cows*. The next step is to define normal and abnormal behavior based on the classified activities which would enable the system to provide information on the individual cow's health and welfare condition.

The severity of false positives and false negatives may vary from one domain to another. For pointing out *focus cows*, it is of high importance to avoid such false classifications as they may lead an animal to suffer from either bad health or welfare conditions without anyone noticing it. As a consequence, it leads to lower production and often an increase of medical expenses. We found, that except for a few values the *standing* distribution tends to collide with the *lying* distribution for all feature distributions. In addition, for the 30% fractile of the *eating seeking* distribution it collides with both *lying* and *standing* distributions. Therefore in a future work we will be searching for features that diverge for the three activities in an attempt to decrease the number of false classifications.

The activity of a cow *drinking* is a useful additional activity to recognize, and it would serve as an important input to recognize abnormal behavior, along with the four activities recognized in this work. By assuming that a cow spending time at a drinking vessel is actually drinking, it would be possible to recognize drinking activity when a cow is in the proximity of a drinking vessel, e.g., by introducing a location model with specific meta information annotated with specific locations. In our future work we plan to include this activity when trying to recognize abnormal behavior.

7 Conclusion

We managed to obtain an average classification success rate of 86.2% for the four activities, by preprocessing position data from cows collected via low-cost GPS receivers, followed by extraction of several features used as input to a standard machine learning technique. The average success rate is higher than we initially expected it to be. The relative high average is reached thanks to the two activities defined as *walking* and *eatingSeeking*, which we recognize in 100.0% and 90.0% of the cases respectively. A challenge for future work lies within the recognition of and distinction between the two activities defined as *standing* and *lying*, where we recognize only 75.8% and 76.5% of the cases respectively. Furthermore, recognizing the transitions between activities will be a future challenge.

References

1. Agouridis, C., Stombaugh, T., Workman, S., Koostra, B., Edwards, D.: Examination of GPS Collar Capabilities and Limitations for Tracking Animal Movement in Grazed Watershed Studies. In: ASAE Annual International Meeting, pp. 27–30 (July 2003)
2. Galileo - a european global navigation satellite system, `http://ec.europa.eu/enterprise/policies/satnav/index_en.htm` (Online; accessed 13-01-2011)
3. Gonzalez, L., Tolkamp, B., Coffey, M., Ferret, A., Kyriazakis, I.: Changes in feeding behavior as possible indicators for the automatic monitoring of health disorders in dairy cows. Journal of Dairy Science 91(3), 1017 (2008)

4. Huzzey, J., Veira, D., Weary, D., von Keyserlingk, M.: Prepartum behavior and dry matter intake identify dairy cows at risk for metritis. Journal of Dairy Science 90(7), 3220–3233 (2007)
5. Ito, K., von Keyserlingk, M., LeBlanc, S., Weary, D.: Lying behavior as an indicator of lameness in dairy cows. Journal of Dairy Science 93(8), 3553–3560 (2010)
6. Kjærgaard, M.B., Blunck, H., Godsk, T., Toftkjær, T., Christensen, D.L., Grønbæk, K.: Indoor positioning using GPS revisited. In: Floréen, P., Krüger, A., Spasojevic, M. (eds.) Pervasive Computing. LNCS, vol. 6030, pp. 38–56. Springer, Heidelberg (2010)
7. Mobile action, `http://www.i-gotu.com/` (Online; accessed 20-12-2010)
8. Mohr, M., Krustrup, P., Bangsbo, J.: Match performance of high-standard soccer players with special reference to development of fatigue. Journal of Sports Sciences 21(7), 519–528 (2003)
9. Nadimi, E., Søgaard, H., Bak, T.: ZigBee-based wireless sensor networks for classifying the behaviour of a herd of animals using classification trees. Biosystems Engineering 100(2), 167–176 (2008)
10. Phillips, C.: Cattle behavior and welfare. Blackwell Science Ltd., Malden (2002)
11. Robert, B., White, B., Renter, D., Larson, R.: Evaluation of three-dimensional accelerometers to monitor and classify behavior patterns in cattle. Computers and Electronics in Agriculture 67(1-2), 80–84 (2009)
12. Schwager, M., Anderson, D., Butler, Z., Rus, D.: Robust classification of animal tracking data. Computers and Electronics in Agriculture 56(1), 46–59 (2007)
13. Spencer, M., Lawrence, S., Rechichi, C., Bishop, D., Dawson, B., Goodman, C.: Time-motion analysis of elite field hockey, with special reference to repeated-sprint activity. Journal of Sports Sciences 22(9), 843–850 (2004)
14. Weka api, the university of waikato, `http://weka.wikispaces.com` (online; accessed 05-01-2011)
15. Zheng, Y., Liu, L., Wang, L., Xie, X.: Learning transportation mode from raw gps data for geographic applications on the web. In: WWW 2008: Proceeding of the 17th International Conference on World Wide Web, Beijing, China, pp. 247–256. ACM, New York (2008)