



HAL
open science

Evaluating ROCKET and Catch22 features for calf behaviour classification from accelerometer data using Machine Learning models

Oshana Dissanayake, Sarah E. Mcpherson, Joseph Allyndrée, Emer Kennedy, Padraig Cunningham, Lucile Riaboff

► To cite this version:

Oshana Dissanayake, Sarah E. Mcpherson, Joseph Allyndrée, Emer Kennedy, Padraig Cunningham, et al.. Evaluating ROCKET and Catch22 features for calf behaviour classification from accelerometer data using Machine Learning models. *Computers and Electronics in Agriculture*, In press. hal-04622072

HAL Id: hal-04622072

<https://hal.science/hal-04622072>

Submitted on 24 Jun 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Evaluating ROCKET and Catch22 features for calf behaviour classification from accelerometer data using Machine Learning models

Oshana Dissanayake^{a,c}, Sarah E. McPherson^{c,d,e}, Joseph Allyndrée^b, Emer Kennedy^{c,d}, Pádraig Cunningham^a, Lucile Riaboff^{a,c,f}

^a*School of Computer Science, University College Dublin, Ireland*

^b*School of Maths and Stats, University College Dublin, Ireland*

^c*VistaMilk SFI Research Centre, Ireland*

^d*Teagasc, Animal & Grassland Research and Innovation Centre, Moorepark, Fermoy, Co. Cork, P61C997, Ireland*

^e*Animal Production Systems Group, Wageningen University & Research, Wageningen, The Netherlands*

^f*GenPhySE, Université de Toulouse, INRAE, ENVT, 31326, Castanet-Tolosan, France*

Abstract

Monitoring calf behaviour continuously would be beneficial to identify routine practices (e.g., weaning, transport, dehorning, etc.) that impact calf welfare in dairy farms. In that regard, accelerometer data collected from neck collars can be used along with Machine Learning models to classify calf behaviour automatically. However, further development is needed to classify a broad spectrum of behaviours with good genericity from one animal to another. While Hand-Crafted features are typically used in the field as inputs for Machine Learning models, feature sets designed explicitly for time-series classification problems have been developed in related fields, such as ROCKET and Catch22 features. This study aims to compare the performance of ROCKET and Catch22 features to Hand-Crafted features commonly used in the field. 30 Irish Holstein Friesian and Jersey pre-weaned calves were equipped with an accelerometer sensor for several weeks, and their behaviours were annotated, allowing for 27.4 hours of observation aligned with the accelerometer time-series. Additional time-series were computed from the raw X, Y and Z-axis and split into 3-second time windows. ROCKET, Catch22 and Hand-Crafted features were calculated for each time window, and the dataset was then split into the train, validation and test sets. Each set of features was used to train three Machine Learning models (Random Forest, eXtreme Gra-

dient Boosting, and RidgeClassifierCV) to classify six behaviours indicative of pre-weaned calf welfare (drinking milk, grooming, lying, running, walking and *other*). Models were tuned with the validation set, and the performance of each feature-model combination was evaluated with the test set. The best performance across the three models was obtained with ROCKET [average balanced accuracy \pm standard deviation] (0.70 ± 0.07), followed by Catch22 (0.69 ± 0.05), well ahead of Hand-Crafted (0.65 ± 0.034). The best balanced accuracy (0.77) was obtained with ROCKET and RidgeClassifierCV, followed by Catch22 and Random Forest (0.73). Thus, tailoring these approaches for specific behaviours and contexts will be crucial in advancing precision livestock farming and enhancing animal welfare on a larger scale.

Keywords: Dairy calf, Behavior, Accelerometers, ROCKET, Catch22, Machine Learning, Features

1. Introduction

Enhancing the welfare of young farm animals through the adoption of suitable practices is likely to improve their performance at different scales. Indeed, the prolonged effects of stress can lead to an exhaustion phase, resulting in decreased performance, an increased risk of disease, and slowed growth. In particular, calves are subjected to many stressful events from their first weeks (dehorning, weaning, transport, social isolation, relocation, etc.). Improving calf welfare is thus highly important to prevent physiological changes and death that may happen due to prolonged exposure to stress (Koknaroglu and Akunal, 2013) but also to bridge the gap between farming and society, as consumers place animal welfare as one of their primary expectations (Cardoso et al., 2016). Changes in calf behaviour, such as altered feeding patterns, decreased playing behaviour or social isolation, can signal underlying health concerns or environmental stressors (Mahendran et al., 2023; Nikkhah and Alimirzaei, 2023). Therefore, monitoring calf behaviour can be a way to identify stress factors (Dissanayake et al., 2022) and to make recommendations on routine practices that are less stressful to promote calf welfare in farming (McPherson et al., 2022). Monitoring the changes in behaviour at an early stage can also help to detect health issues or discomfort as soon as they arise to ensure rapid intervention by the farmer, thus reducing the use of medications and their associated costs. However, automatic and continuous behaviour monitoring is required for the targeted applications,

which is not compatible with human observations that are time-consuming and labour-intensive (Penning, 1983). In that regard, sensors that can monitor livestock behaviour automatically have been increasingly used over the last few years (Rushen et al., 2012; Jiang et al., 2023). In particular, onboard accelerometer sensors offer a continuous data stream in the three-dimensional plane from which a diverse range of movements and activities can be derived. These sensors are also adaptable and versatile, can function for a very long period, and can be integrated into various devices and systems, explaining their rapid adoption in livestock (Martiskainen et al., 2009; Moreau et al., 2009; Tatler et al., 2018; Chakravarty et al., 2019; Yu et al., 2023). For example, González et al. (2015) have shown that ruminating can be classified with a sensitivity $> 85\%$ and specificity $> 90\%$ based on a decision tree classifier. Iqbal et al. (2021) used a decision tree algorithm to classify data segments into foraging, ruminating, travelling, resting, or *other* behaviours, with an accuracy of $> 85\%$. Similarly, Benaissa et al. (2019) have shown that the precision and sensitivity for lying behaviour were above 93% using a leg-mounted accelerometer sensor with Machine Learning models.

However, Riaboff et al. (2022) highlighted in a comprehensive review that the challenge lies in capturing a large spectrum of behaviours, including those only expressed occasionally. For instance, transitional behaviours such as lying-down and standing-up are not frequently observed and are often poorly predicted, with a sensitivity lower than 70% (Martiskainen et al., 2009; Vázquez Diosdado et al., 2015). Some maintenance behaviours (e.g., urinating, drinking), self-grooming behaviours and social interaction also tend to get accuracies lower than 80% (Lush et al., 2018; Rodríguez-Baena et al., 2020). Recently, Hosseininoorbin et al. (2021) stated again that cattle behaviour classification from accelerometer data can provide excellent performance for a small number of classes, but the performance decreases substantially when more than 5 behaviours are considered. However, both core time-budget behaviours and less frequent behaviours are essential for assessing calf welfare. Indeed, the time spent lying-down is usually increased around the time of the diagnosis of respiratory diseases (Duthie et al., 2021). Enríquez et al. (2010) observed that playing stopped abruptly after the physical separation from the dams, while a peak of seeking and walking behaviours was observed the day after the separation. Fewer self-grooming and more scratching events are recorded in the few hours after dehorning, while more lying-down and standing-up transitions are observed (Morisse et al., 1995). Furthermore, Riaboff et al. (2022) also highlighted concerns

regarding the generic nature of the models. While the model performance is usually high when the model is tested with data from animals already used for model training, a significant drop in performance is observed when the model is tested with new animals, due to a lack of similarity between the feature spaces of the training and test sets (Rahman et al., 2018). In calves, an extensive range of behaviours, including locomotor-play, self-grooming, ruminating, non-nutritive suckling, nutritive suckling, active lying and non-active lying, has been predicted with an accuracy higher than 90% using the AdaBoost algorithm. However, all the behavioural observations and raw sensor data from the subjects have been merged according to timestamps to create the datasets (Carslake et al., 2020). Thus, though cross-validation has been performed in model testing, the genericity of the model over different calves has yet to be assessed. This limitation compromises the deployment on experimental and commercial farms, as the models must be robust regardless of the animal. Therefore, improving (i) the range of behaviours correctly predicted and (ii) the genericity of the models are both necessary to evaluate animal welfare and health status, particularly in calves, which have received less attention than cows despite the issues surrounding the welfare of young livestock.

The prevailing methodology employed in literature for animal behaviour classification entails the segmentation of the accelerometer signal in fixed-size time windows (Riaboff et al., 2022). Hand-Crafted (**HC**) features are then calculated in each time window to capture the dynamic properties of the time-series, predominantly rooted in the time and frequency domains, encompassing mean, median, motion variation, roll, pitch, and spectral entropy (Lush et al., 2018; Chang et al., 2022; Price et al., 2022; Dissanayake et al., 2022; Balasso et al., 2023). These features are then used to train a Machine Learning model, usually RandomForest (**RF**), Linear Discriminant Analysis, Support Vector Machines, Deep Neural Networks, or eXtreme Gradient Boosting. (Riaboff et al., 2020; Kleanthous et al., 2022). While the performance of a set of models has often been compared (Dutta et al., 2015; Smith et al., 2016; Hu et al., 2020), the impact of the feature sets used to classify behaviours in livestock ruminant has yet to be explored. On the contrary, extensive studies in the time-series domain have already highlighted some sets of features that perform very well for classification problems in different applications. In particular, Lubba et al. (2019) have identified a set of 22 non-correlated CAnonical Time-series CHaracteristics (called Catch22) initially derived from a large pool of 4791 features based on their exceptional per-

formance across 93 time-series classification datasets. The Catch22 features include the most efficient features, covering auto-correlation, value distributions, outliers, and fluctuation scaling properties, while minimizing redundancy. Moreover, Dempster et al. (2020, 2021) recently introduced Random Convolutional Kernel Transform (ROCKET) features that surpasses conventional approaches in time-series classification problems. ROCKET uses Random Convolutional Kernels, with kernels of random length, weights, bias, dilation and padding to transform raw time-series into high-dimensional feature representations. In that way, ROCKET captures diverse temporal patterns without domain-specific knowledge to distinguish underlying structures and patterns within the data. Although the importance of individual features within different HC feature sets is frequently documented (Lush et al., 2018; Dissanayake et al., 2022), the performance of a set of features developed explicitly for time-series classification problems in related fields, such as Catch22 and ROCKET, has never been investigated. Therefore, this study aims to evaluate the effectiveness of Catch22 and ROCKET for classifying calf behaviour from accelerometer data. Especially, we evaluate Catch22 and ROCKET features performance compared to HC features considering the two main limitations highlighted in the literature, i.e., (i) discriminating a large spectrum of pre-weaned calf behaviours, including drinking milk, running, lying, walking, grooming and *other* and (ii) maintaining good performance when the model is applied to new calves than those used for model training.

2. Materials and Methods

An overview of the methodology applied is described in Figure 1. Each step is then detailed in the following sections.

2.1. Data Collection

The experiment was carried out at Teagasc Moorepark Research Farm (Fermoy, Co. Cork, Ireland; $50^{\circ}07'N$; $8^{\circ}16'W$) from January 21 to April 5, 2022. Ethical approval was obtained from the Teagasc Animal Ethics Committee (TAEC; TAEC2021–319). The trial was carried out in accordance with the European Union (Protection of Animals Used for Scientific Purpose) Regulations 2012 (S.I. No. 543 of 2012). 47 Irish Holstein Friesian and Jersey pre-weaned calves were utilized for the experiment. The calves were managed according to conventional rearing and management practices (Conneely et al., 2014) at Teagasc Moorepark Research Farm. After calving, the calves and

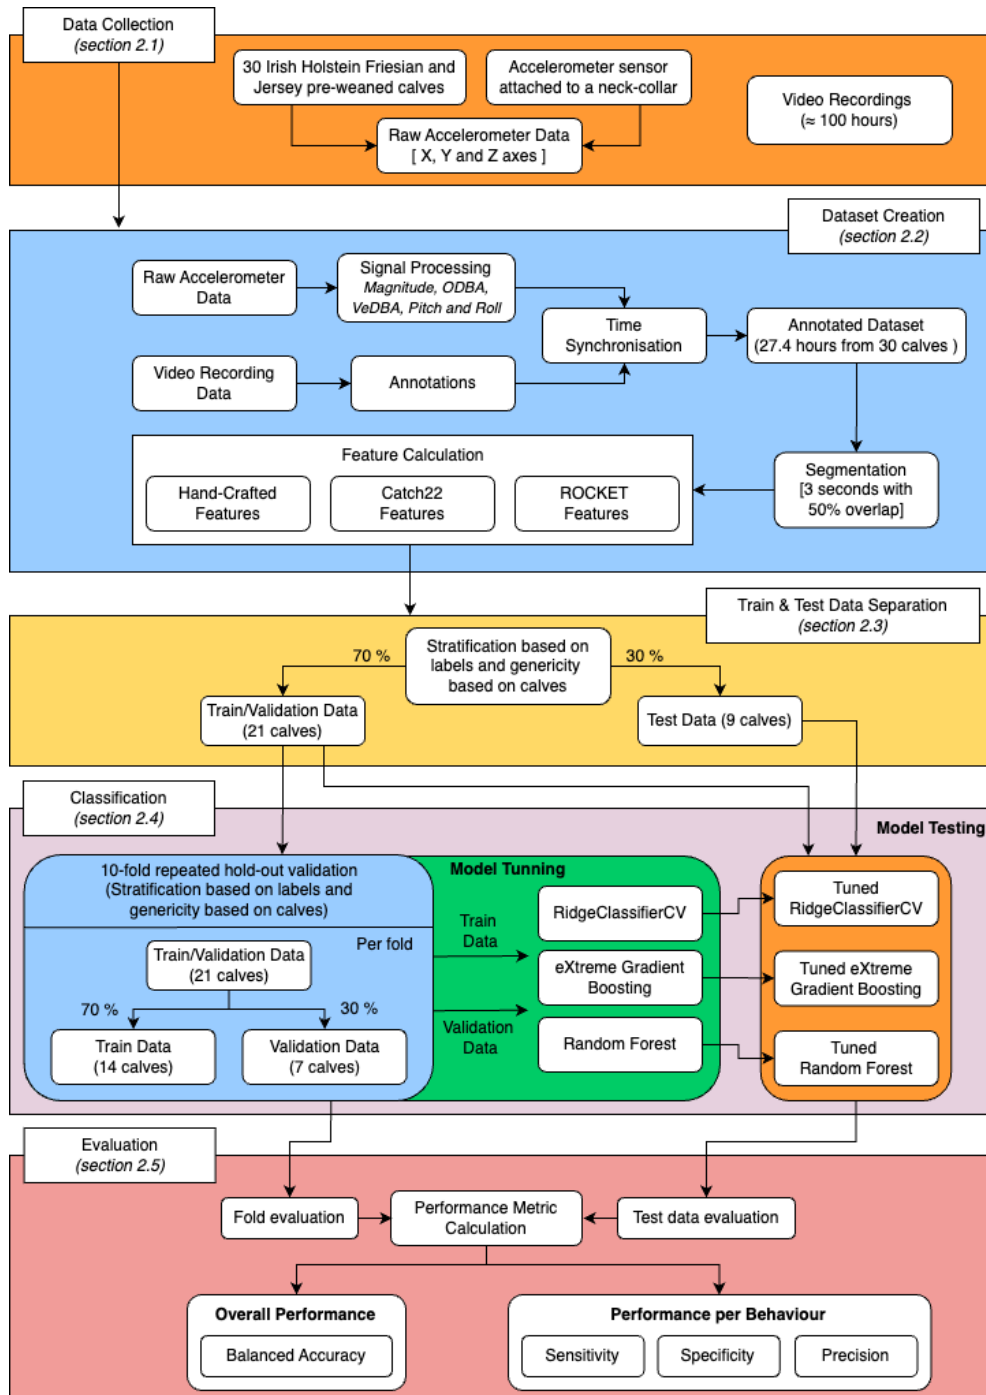


Figure 1: Overview of the methodology applied to compare classification performance depending on the subset of features (HC, Catch22 and ROCKET) and models (RF, XGB, RidgeClassifierCV) and evaluate model genericity across calves.

dams were separated within one hour. The calves were moved to a straw-bedded individual pen and were artificially fed with their mother’s colostrum $< 2h$ post-birth at a rate of 8.5% of their birth weight. After receiving colostrum, calves were fed their own dam’s transition milk at a rate of 10% of their birth weight twice a day for their subsequent five feedings. After transition milk, calves were fed a 2.5L milk replacer (26% crude protein; Volac Heiferlac Instant, Volac, Hertfordshire, UK) twice daily. Between 3-7 days old, the calves were transferred to a group pen (see Figure 2), where they were fed using an automatic milk feeder at a rate of 6 L/calf/day with *ad libitum* access to hay, concentrates, and water. Calves were gradually weaned at 56 days using the automatic feeder.




Figure 2: Design of the group pens used after 3-7 days old of age with cameras placement to film the calves in each pen.

Each calf was equipped with a tri-axis accelerometer data logger (Axivity LTD¹) fastened to a neck collar starting from one week after birth until 2 months of age (see Figure 3). The accelerometers were configured with a sensitivity of $\pm 8g$ and a sampling rate of 25 Hz (battery life: 30 days). A NAND flash memory was used to store the data (memory: 512Mb). The accelerometers were measuring $23 \times 32.5 \times 7.6$ mm and weighing 11g. Each sensor was wrapped in cling film and cotton wool and attached to the collar with vet wrap and insulating tape. The sensors were placed on the left side of the neck in the same orientation for all calves (see Figure 3). The X-axis detected the top-bottom direction, the Y-axis detected the backwards-forward direction, and the Z-axis detected the left-right direction. The collars

¹Axivity Ltd; <https://axivity.com/product/ax3>




Figure 3: The collar used to attach the accelerometer sensor and the orientation of the three axes. The X-axis detected the top-bottom direction, the Y-axis detected the backwards-forward direction, and the Z-axis detected the left-right direction.

were tightly adjusted, and a 13g metal ring was added to prevent them from moving from the designated side (Figure 3). For the next ten weeks, collars were taken off every two weeks to retrieve data and replenish the battery. A set of DVRs (4-channel and 8-channel Hikvision² 1080p) and four 8Mp Dome³ CCTV cameras mounted in each pen were employed in addition to the accelerometer data collection to record videos of calves (see Figure 2). 2092 hours of videos were extracted at the end of the trial.

A subset of 30 calves out of the 47 who demonstrated a wide range of behaviours within the selected videos were selected to create the dataset for this study. Roughly 3 hours of video time were inspected per calf out of multiple videos. This was done to ensure that the selected videos included a wide range of behaviours displayed by each calf across different situations. Labelling was done with the Behavioural Observation Research Interactive Software (BORIS) (Friard and Gamba, 2016) using an exhaustive ethogram with 24 behaviours based on Barry et al. (2019) after adaptation to this experiment. Annotations were carried out by three observers, ensuring that all 30 animals were observed for at least 15 minutes by at least one observer. The three observers annotated a one-hour video for one calf to measure the concordance between annotations every second based on the 24 behaviours (Cohen’s Kappa averaged over the 1-hour video: $0.72 + / - 0.01$). A total

²Equicom Ltd; <https://www.equicom.ie/8-channel-dvr-recorder-for-cctv-cameras.html>

³Equicom Ltd; <https://www.equicom.ie/1080p-mini-dome-cctv-camera.html>

of 27.4 hours of observation has been performed over the 30 calves (age: 23.7 ± 10.7 days). Finally, five behaviours indicative of calf welfare each representing more than 1.7% of the whole dataset were kept in the study, i.e., lying (42.1%), drinking milk (9.23%), grooming (4.29%), running (2.24%) and walking (1.71%). All the other behaviours were merged into a class *other* (40.42%) and retained in the rest of the process. The definitions of each behaviour are listed in Table 1 and, the amount of data collected for each behaviour is detailed in Table 2.

The accelerometer time-series were aligned with the observations after synchronizing the accelerometer timestamps with the video timestamps. This was achieved using an external reference clock used during the data collection to make the correspondence between the accelerometer and camera timestamps.

Table 1: Definition of the different behaviours (Barry et al., 2019) in the ethogram used for that study after merging the least observed behaviours (<1.7 %) in the *other* class.

Behaviour	Definition
Drinking milk	Calf is drinking milk from the milk feeder.
Grooming	Calf uses tongue to repeatedly lick own back, side, leg, tail areas.
Lying	Calf is resting either sternally or laterally with all four legs hunched close to body either awake or asleep.
Running	Calf is running (play / not-play).
Walking	Calf is walking or shuffling about.
<i>Other</i>	A collection of 19 other behaviours including rising, lying-down, social interaction, play etc.

2.2. Creation of the 3 features datasets

The computational workflow was executed on an Intel Xeon E-2378G CPU (2.80GHz and based on x86_64 architecture) with 16 CPUs, eight cores per socket, and two threads per core with a Matrox G200eW3 GPU.

2.2.1. Accelerometer time-series calculation

A set of time-series typically calculated in the literature was derived from the X, Y, and Z axes, as recommended in Riaboff et al. (2022). The magnitude of the acceleration was first computed from the raw X, Y and Z axes

Table 2: Summary of data collected: Duration of observation, number of sequences associated with each behaviour, and number of calves on which each behaviour was observed.

Behaviour	Total Duration (minutes)	Number of sequences	Number of calves
Drinking milk	273.6	169	27
Grooming	127.35	334	29
Lying	1248.35	120	27
Running	66.45	608	24
Walking	50.75	561	30
<i>Other</i>	1198.6	2636	30

readings of the accelerometer, removing the gravity component of the acceleration (van Hees et al., 2013; da Silva et al., 2014) (equation 1). A 6th-order Butterworth high-pass filter with a cut-off frequency of 0.3 Hz (Smith et al., 2016) was applied to isolate the acceleration component linked to the movement of the animal’s body (dynamic acceleration). Overall Dynamic Body Acceleration (ODBA) (Wilson et al., 2008; Versluijs et al., 2023) and Vectorial Dynamic Body Acceleration (VeDBA) (Walker et al., 2015; Versluijs et al., 2023) were computed from the dynamic acceleration according to the equations 2 and 3, respectively. A 6th-order Butterworth low-pass filter with a cut-off frequency of 0.3 Hz was applied to isolate the acceleration component linked to gravity (static acceleration). Pitch (Walker et al., 2015; Versluijs et al., 2023) and roll (Walker et al., 2015; Versluijs et al., 2023) were computed from the static component according to equations 4 and 5, respectively.

$$magnitude = \sqrt{(A_x^2 + A_y^2 + A_z^2)} - 1g \quad (1)$$

$$ODBA = |D_x| + |D_y| + |D_z| \quad (2)$$

$$VeDBA = \sqrt{(D_x^2 + D_y^2 + D_z^2)} \quad (3)$$

$$pitch = \arctan(S_z / (S_y^2 + S_x^2)) \quad (4)$$

$$roll = \arctan(S_y / (S_z^2 + S_x^2)) \quad (5)$$

where A_x, A_y and A_z represents the X, Y and Z axes reading of the accelerometer, D_x, D_y , and D_z represent the dynamic component of the acceleration and S_x, S_y and S_z the static component of the acceleration.

2.2.2. Segmentation into fixed time-windows

Each time-series was split into a 3 seconds-window to (i) meet the [3 - 30 seconds] range as recommended in the literature to get enough signal information within a window (Riaboff et al., 2022) and (ii) get a time-window short enough to adapt to the brief and changing behaviours of pre-weaned calves. As the literature recommends, a 50% overlap has been applied between each window (Riaboff et al., 2022). It should be noted that the best window size and overlap combination was not investigated in this study as our main objective was to compare the performance of different subsets of features along with Machine Learning models rather than optimising a methodological framework.

2.2.3. Feature Calculation

ROCKET, Catch22, and Hand-Crafted (HC) features were calculated from each 3-seconds window across the 8 time-series: X, Y, and Z axes, magnitude, OBDA, VeDBA, pitch and roll.

1. Hand-Crafted features

HC features calculated in livestock ruminant classification behaviour from accelerometer data are usually extracted in the time domain and frequency domain using the Fourier Transform. Those features provide information on the motion intensity (e.g., median, quartiles, maximum, motion variation from ODBA and VeDBA), the orientation of the animal’s body (e.g., pitch and roll), the shape of the signal distribution (e.g., skewness and kurtosis) and the physical properties of the movement (periodic, stochastic, etc.) (e.g., spectral entropy, fundamental frequency, etc.) (Riaboff et al., 2022). In this study, we calculated mean, median, minimum, maximum, standard deviation, first quartile, third quartile, skewness, kurtosis, entropy and motion variation. Across the 8 time-series, this led to a set of 88 HC features typically calculated in the field, providing information on the four categories listed above. This feature set was used as the reference to evaluate the potential additional value of Catch22 and ROCKET. Features, extraction domain, equations and references are displayed in Table 3.

Table 3: Hand-Craft features calculated as the reference: Equations, domain and references.

Feature	Domain	Equation	Reference
Mean	Time	Mean = $\frac{\sum_1^N X_i}{N}$ N = Total number of observations.	(Preece et al., 2008) (Drover et al., 2017) (Tatler et al., 2018)
Median	Time	$M = \begin{cases} X_{\frac{n+1}{2}}, & \text{if } n \text{ is odd} \\ \frac{X_{\frac{n}{2}} + X_{\frac{n}{2}+1}}{2}, & \text{if } n \text{ is even} \end{cases}$ X = Data points in the sorted list.	(Preece et al., 2008) (Figo et al., 2010) (Fida et al., 2015)
Minimum	Time	Min(X) = $\min(X_1, X_2, \dots, X_n)$ X = Considered set of data points.	(Figo et al., 2010) (Barwick et al., 2018)
Maximum	Time	Min(X) = $\max(X_1, X_2, \dots, X_n)$ X = Considered set of data points.	(Figo et al., 2010) (Barwick et al., 2018)
Standard Deviation	Time	$\sigma = \sqrt{\frac{\sum(x_i - \mu)^2}{N}}$ N = size of the data segment. x_i = Individual observation in the data set. μ = mean of the data segment.	(Preece et al., 2008) (Figo et al., 2010) (Bersch et al., 2014) (Drover et al., 2017)
First Quartile	Time	$Q1 = \frac{N+1}{4}$ N = Total number of observations in the data set.	(Preece et al., 2008) (Zdravetski et al., 2017) (Fan et al., 2019)
Third Quartile	Time	$Q3 = \frac{3}{4}(N + 1)$ N = Total number of observations in the data set.	(Preece et al., 2008) (Zdravetski et al., 2017) (Fan et al., 2019)
Spectral Entropy	Frequency	$H(s, sf) = \sum_{f=0}^{f_s/2} P(f) \log_2[P(f)]$ Where P is the normalised power spectral density, and f_s is the sampling frequency.	(Preece et al., 2008) (Riaboff et al., 2020) (Aziz et al., 2021) (Dissanayake et al., 2022)
Motion Variation	Time	$MV = \frac{1}{M} \left(\sum_{i=1}^{M-1} a_{x,i+1} - a_{x,i} + \sum_{i=1}^{M-1} a_{y,i+1} - a_{y,i} + \sum_{i=1}^{M-1} a_{z,i+1} - a_{z,i} \right)$	(Riaboff et al., 2020) (Fogarty et al., 2020) (Dissanayake et al., 2022)
Skewness	Time	$\gamma = \frac{1}{N} \sum_{i=1}^N \left(\frac{Y_i - \mu}{\sigma} \right)^3$ γ = skewness N = number of variables in the distribution μ = mean of the distribution σ = standard deviation	(AlZubi et al., 2014) (Hounslow et al., 2019) (Cabezas et al., 2022)
Kurtosis	Time	$\beta = \frac{\frac{1}{N} \sum_{i=1}^N (Y_i - \mu)^4}{V^2}$ β = kurtosis N = number of variables in the distribution μ = mean of the distribution V = variance of the dataset	(AlZubi et al., 2014) (Hounslow et al., 2019) (Cabezas et al., 2022)

2. Catch22 features

Catch22 are 22 Canonical Time-series Characteristics derived from a large pool of 4791 features that showed exceptional performance across 93 time-series classification datasets. This subset was also selected to minimize redundancy between features, thus providing complementary information on the time-series while reducing computational expenses (Lubba et al., 2019). The Catch22 features cover auto-correlation, value distributions, outliers, and fluctuation scaling properties. This set can be increased to 24 features by including mean and standard deviation to include the location and spread of the raw time-series distribution in the classification process. As that information may be necessary for the time-windows classification into calf behaviour, the set of 24 features was calculated, leading to a subset of 192 features across the 8 time-series. For the remainder of the paper, we are keeping the name *Catch22* as soon as we refer to that features set. Feature names and their description are displayed in Table 4.

Table 4: Catch22 features for time-series classification problem (Lubba et al., 2019).

Feature name	Description
DN_HistogramMode_5	Mode of z-scored distribution (5-bin histogram)
DN_HistogramMode_10	Mode of z-scored distribution (10-bin histogram)
SB_BinaryStats_mean_longstretch1	Longest period of consecutive values above the mean
DN_OutlierInclude.p_001_mdrmd	Time intervals between successive extreme events above the mean
DN_OutlierInclude.n_001_mdrmd	Time intervals between successive extreme events below the mean
CO_1/ecac	First $1/e$ crossing of autocorrelation function
CO_FirstMin_ac	First minimum of autocorrelation function
SP_Summaries_welch_rect_area_5_1	Total power in lowest fifth of frequencies in the Fourier power spectrum
SP_Summaries_welch_rect_centroid	Centroid of the Fourier power spectrum
FC_LocalSimple_mean3_stderr	Mean error from a rolling 3-sample mean forecasting
CO_trev_1_num	Time-reversibility statistic, $(xt + 1 - xt)3t$
CO_HistogramAMI_even_2_5	Automutual information, $m = 2, \tau = 5$
IN_AutoMutualInfoStats_40_gaussian_fmml	First minimum of the automutual information function
MD_hrv_classic_pnn40	Proportion of successive differences exceeding 0.04σ (Mietus et al., 2002)
SB_BinaryStats_diff_longstretch0	Longest period of successive incremental decreases
SB_MotifThree_quantile_hh	Shannon entropy of two successive letters in equiprobable 3-letter symbolization
FC_LocalSimple_mean1_ttauresrat	Change in correlation length after iterative differencing
CO_Embed2_Dist_tau_d_expfit_meandiff	Exponential fit to successive distances in 2-d embedding space
SC_FluctAnal_2_dfa_50_1_2_logi_prop_r1	Proportion of slower timescale fluctuations that scale with DFA (50% sampling)
SC_FluctAnal_2_rsrangeft_50_1_logi_prop_r1	Proportion of slower timescale fluctuations that scale with linearly rescaled range fits
SB_TransitionMatrix_3ac_sumdiagcov	Trace of covariance of transition matrix between symbols in 3-letter alphabet
PD_PeriodicityWang_th0.01	Periodicity measure of (Wang et al., 2007)

3. ROCKET features

ROCKET (RandOm Convolutional KErnel Transform) is a pioneering technique in time-series classification, known for its efficiency and efficacy in features extraction. ROCKET uses Random Convolutional

Kernels, where kernels have random lengths, weights, biases, dilations, and paddings, to transform raw time-series into high-dimensional feature representations. ROCKET features are specifically the maximum and proportion of positive values resulting from each convolution, meaning that for k kernels, $2k$ features per time-series are produced. Using random kernels allows to capture diverse temporal patterns without domain-specific knowledge. For this study, we have utilized *miniROCKET* which leverages a more streamlined set of convolutional kernels, focusing on a predefined subset of kernel parameters rather than the exhaustive randomness of ROCKET. This refinement allows miniROCKET to achieve similar or even superior classification performance with a significantly reduced computational footprint (Dempster et al., 2021). The default number of kernels (10000) was used during the feature generation in that study. ROCKET features were calculated from each of the 8 time-series, leading to a set of 9996 features.

2.3. Partitioning the dataset into training and testing calf-independent sets

The dataset was split into a 70:30 calf ratio (Figure 1): Out of 30 calves, 21 were chosen for the training set, and 9 were chosen for the test set. This split ensures that the calves used for testing the model have not been used for model training to evaluate the model genericity. Furthermore, stratification has been applied to the annotated behaviours to maintain a consistent proportional distribution between the train and test set. The optimal split used in the rest of the study is the one that minimises the mean of the differences in the proportions of each of the annotated behaviours between the training and testing sets.

2.4. Modelling with Machine Learning models

Three Machine Learning models were used along with each feature set to (i) conclude on the average performance of the feature sets across several models and (ii) investigate whether some feature set-model combinations perform better than others. eXtreme Gradient Boosting (**XGB**), RandomForest (**RF**) and RidgeClassifierCV (**RCV**) algorithms were used, considering their high performance in classification problems in various domains.

1. eXtreme Gradient Boosting

The eXtreme Gradient Boosting (XGB) algorithm is an ensemble Ma-

chine Learning technique that uses gradient boosting techniques to improve model accuracy (Friedman, 2001). It constructs decision trees sequentially, with each tree correcting previous errors to enhance model accuracy. The algorithm uses gradient boosting, where each tree is trained using the gradient of the loss function, minimizing the difference between predicted and actual values. XGB is highly versatile, capable of handling sparse data and working with various data formats. It also incorporates techniques to prevent overfitting, such as regularization terms in the objective function. XGB offers efficient scalability and parallel processing capabilities, significantly speeding up computation time, especially with large datasets. Below are a few most prominent hyper-parameters (Chen and Guestrin, 2016). The values tested for each hyperparameter in this study are shown in Table 5.

- **eta (learning rate):** Determines the step size at each iteration while moving toward a minimum of the loss function. A lower value makes the model more robust at the cost of slower computation.
- **max depth:** Sets the maximum depth of a tree. Increasing this value will make the model more complex and likely to overfit.
- **n_estimators:** Number of trees to be built. More trees can increase accuracy but also computation time.
- **gamma:** Minimum loss reduction required to make a further partition on a leaf node of the tree. The larger the value, the more conservative the algorithm will be.
- **lambda (reg lambda) and alpha (reg alpha):** These are L2 and L1 regularization terms on weights, respectively, and can be used to handle overfitting.
- **scale_pos_weight:** Controls the balance of positive and negative weights, useful for unbalanced classes.

2. Random Forest

The Random Forest algorithm is a versatile ensemble learning method for classification tasks. It constructs multiple decision trees during training and outputs the mode of classes or mean prediction of the individual trees. Randomness is introduced through bootstrap sampling and a random subset of features at each node, making the model

more robust and preventing overfitting. The algorithm can handle large datasets effectively with higher dimensionality, performs well with categorical and continuous variables, and can handle missing values (Breiman, 2001). Below are a few most prominent hyper-parameters based on (Pedregosa et al., 2011) and the tested values are shown in Table 5.

- **n_estimators:** The number of trees in the forest.
- **criterion:** The function to measure the quality of a split.
- **max depth:** The maximum depth of the tree.
- **min samples split:** The minimum number of samples required to split an internal node. Higher values prevent creating trees that are too complex and overfitting.
- **max features:** The number of features to consider when looking for the best split. It helps in adding randomness to the model.
- **class weight:** Weights associated with classes. It is particularly useful in scenarios where classes are imbalanced, as it can influence the decision trees in the forest towards the minority class.

3. RidgeClassifierCV

RidgeClassifierCV is useful in large-variable scenarios and aims to prevent overfitting while maintaining a balance between bias and variance. It is efficient for high-dimensional datasets and is commonly used in problems requiring interpretability and prediction accuracy. RidgeClassifierCV (Pedregosa et al., 2011) is an extension of the RidgeClassifier based on the theory of Ridge Regression (Hoerl and Kennard, 1970), which is an extension of linear regression that uses cross-validation to determine the optimal regularization parameter. This method splits the dataset into subsets and evaluates the model's performance for different regularization parameter (alpha) values. The **CV** stands for cross-validation, which helps to find the alpha that yields the best generalization performance. Following are a few important hyperparameters of RidgeClassifierCV (Pedregosa et al., 2011) and the tested values are shown in Table 5.

- **alphas:** Determines the regularization strength. Regularization improves the conditioning of the problem and reduces the variance of the estimates. Larger values promotes stronger regularization.

- **fit_intercept:** Determines whether a bias (intercept) term is added to the decision function, allowing the model to fit data that is not centered around zero.
- **cv:** Determines the cross-validation splitting strategy.
- **class_weight:** Weights associated with classes. It is particularly useful in scenarios where classes are imbalanced.

Table 5: Tested Classifiers and hyper-parameters.

Classifier	Hyperparameters	Values Tested
XGBoost	n_estimators	100, 200
	eta	0, 0.5, 1
	gamma	0, 5, 10
	max_depth	None, 0, 10
	class_weight	None, balanced
Random Forest	n_estimators	100, 200
	max_depth	None, 10
	min_samples_split	2, 5
	max_features	None, log2, sqrt
	criterion	gini, entropy
RidgeClassifierCV	class_weight	None, balanced
	fit_intercept	True, False
	alphas	np.logspace(-3,3,100), np.logspace(-1,10,100)

2.5. Model tuning and evaluation for each feature set-model combination

The Machine Learning models were developed using scikit-learn (*v.1.2.2*) (Pedregosa et al., 2011), XGBoost (*v.1.6.1*) (Chen and Guestrin, 2016), and skTime (*v.0.24.0*) (Löning et al., 2019) available under Python *v.3.9.7*. For each of the 3 sets of features, the 3 Machine Learning models were tuned using the training data. A grid search was used to identify the best hyper-parameters for each model along with each set of features (see Table 6). For that purpose, each model was trained with one of the combinations of hyperparameters in the grid search using 14 calves from the 21 available in the training dataset. The model was then evaluated using a validation dataset made up of the 7 remaining calves in the training set. The process has been iterated 10 times (see Figure 1). For each feature set and algorithm, the model that achieved the best balanced accuracy (BA; see equation 6) on

average over the 10 iterations on the validation dataset was selected for the rest of the study (tuned model). For each set of features, the performances of each of the 3 tuned Machine Learning models were calculated using the BA as a global metric of performance and using the sensitivity, specificity and precision as a performance metric per behaviour (see equation 7, 8 and 9, respectively).

Table 6: Best performing hyper-parameters identified through the grid search for each feature set-model combination.

Model	Hyper-parameters	Hand-Crafted	Catch22	ROCKET
XGB	class_weight	None	None	None
	eta	0.5	1	1
	gamma	0	0	0
	max_depth	10	6	6
	n_estimators	200	200	200
RF	class_weight	balanced	balanced	balanced
	criterion	gini	entropy	gini
	max_depth	10	sqrt	sqrt
	max_features	log2	10	10
	min_samples_split	5	5	2
	n_estimators	200	200	200
RidgeClassifierCV	class_weight	balanced	balanced	balanced
	fit_intercept	False	True	False
	alphas	np.logspace(-1,10,100)	np.logspace(-1,10,100)	np.logspace(-1,10,100)

$$\text{Balanced Accuracy} = \frac{1}{2} \left(\frac{TP}{TP + FN} + \frac{TN}{TN + FP} \right) \quad (6)$$

$$\text{Sensitivity} = \frac{TP}{TP + FN} \quad (7)$$

$$\text{Specificity} = \frac{TN}{TN + FP} \quad (8)$$

$$\text{Precision} = \frac{TP}{TP + FP} \quad (9)$$

Where **TP** (True Positive) is the number of time-windows where the behaviour of interest was observed and correctly predicted; **FN** (False Negative) is the number of time-windows where the behaviour of interest was observed but another behaviour was predicted, **FP** (False Positive) is the number of time-windows where the behaviour of interest was predicted but another be-

behaviour was observed; **TN** (True Negative) is the number of time-windows where the behaviour of interest was not observed and not predicted.

The Python code utilized for the sections 2.2, 2.3, 2.4 and 2.5 is presented through an *open-source Github project* ⁴.

3. Results

3.1. Model tuning with 10-fold cross-validation

The best performing hyperparameters identified through the grid search are detailed in Table 6, and their respective scores achieved during each 10 fold of the cross-validation for Hand-Crafted, Catch22 and ROCKET feature sets are displayed in Figure 4.

On average, highest BA across the 10 folds were obtained with ROCKET (0.70 ± 0.07 (*mean \pm standard-deviation (std)*)) followed by Catch22 (0.69 ± 0.05) and HC features (0.65 ± 0.03) which showed the lowest efficacy.

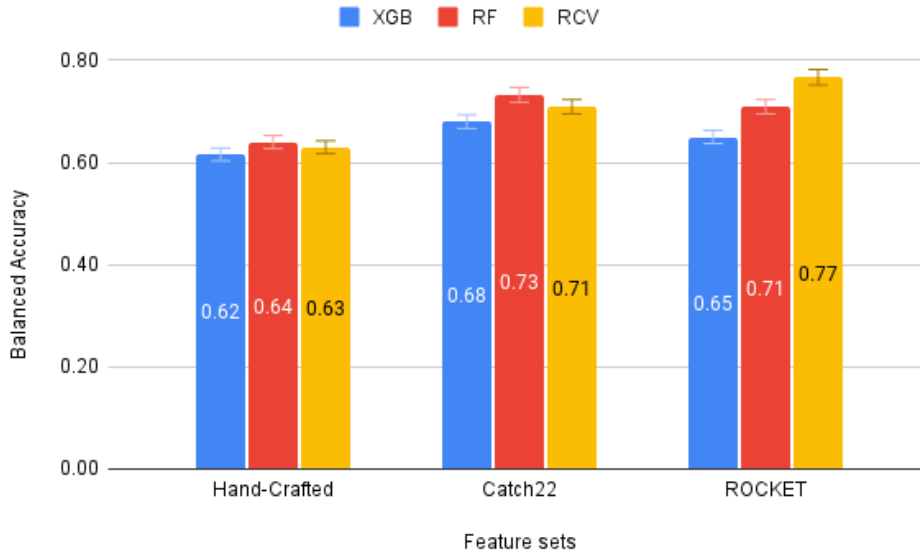


Figure 4: Mean BA with standard-deviation obtained for each 10-folds for every feature set-model combination.

⁴https://github.com/Oshana/comp_n_ele_in_agriculture.git

3.2. Model performance with the test set

3.2.1. Overall Performance

The BA scores obtained for each fitted model and set of features on the calf-independent test set are presented in Figure 5. Overall, the best BA across the 3 Machine Learning models was reached with ROCKET (0.70 ± 0.07 [*mean* \pm *std*]), followed closely by Catch22 (0.69 ± 0.05). BA was substantially lower with HC features (0.65 ± 0.03).

Variability within the same subset of features has been found depending on the Machine Learning model (Figure 5). Indeed, ROCKET obtained its highest and lowest BA with RCV and XGB, respectively. Catch22 and HC features reached their best BA with RF and their lowest BA with XGB.

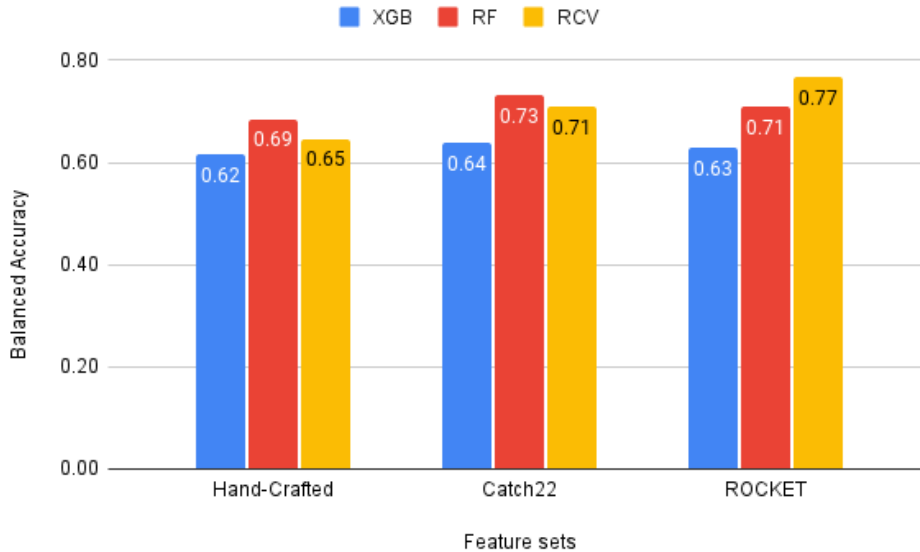


Figure 5: Balanced Accuracy obtained on the calf-independent test-set for every feature set-model combination.

3.2.2. Performance per behaviour

The best sensitivity, specificity and precision for each behaviour across the 3 Machine Learning models were obtained either with ROCKET or Catch22 features, except for the sensitivity of walking (0.65 ± 0.23 [*mean* \pm *std*]), the specificity of *other* (0.83 ± 0.11) and the precision of running (0.93 ± 0.03).

where the best performance were slightly better with HC features (Table 7; see Appendix A for the complete report of performance).

Figure 6 displays the sensitivity, specificity and precision obtained for each behaviour for each feature set-model combination. Lying and running were correctly predicted regardless of the feature set-model combination, with sensitivity, specificity and precision higher than 0.80. However, some feature set-model combinations are more promising for classifying drinking milk, grooming, walking and *other*. As displayed in Table 7, drinking milk was better predicted with ROCKET, especially with RCV and XGB. Similarly, grooming achieved its best performance with ROCKET in combination with RCV and XGB and Catch22 in association with XGB. Walking behaviour was best predicted with HC and Catch22, combined with RCV and XGB. The best performance for predicting *other* behaviour class was achieved by combining Catch22 with XGB and RF, ROCKET with XGB, and HC features with RCV.

In the same way as for the BA, the performances per behaviour are highly variable within each set of features, depending on the Machine Learning model (see Table 7 and Appendix A for the full report). The impact of the model on the feature’s performance is particularly noticeable for the sensitivity and precision of the behaviours of drinking milk, grooming, walking and *other*. For the HC features, a gap in sensitivity of 0.68, 0.45 and 0.30, respectively, has been found between the highest and the lowest performing model combination for the behaviours *other*, walking and drinking milk, respectively. A difference in precision of 0.34, 0.31 and 0.47 has been obtained between the best and worst model combinations with HC for drinking milk, grooming and walking. For the Catch22 features, substantial differences were also observed depending on the model, with a gap in sensitivity of 0.47, 0.26 and 0.37 for the behaviours *other*, drinking milk and walking, respectively. A high difference in the precision of 0.58, 0.35 and 0.26 for walking, grooming and drinking milk has also been obtained between the best and worst model associated with Catch22. For ROCKET features, a difference in sensitivity of 0.45, 0.33 and 0.26 was found for walking, grooming and drinking milk, respectively, depending on the model. At the same time, a gap in precision of 0.44 and 0.40 was recorded between the best and worst model combinations for the behaviours of walking and grooming.

Table 7: Best sensitivity, specificity and precision achieved per feature set-model combination per behavior.

	HC			Catch22			ROCKET		
Machine learning model	XGB	RF	RCV	XGB	RF	RCV	XGB	RF	RCV
Best sensitivity									
Drinking milk			0.71						0.71
Grooming									0.60
Lying						0.90			
Running						0.99			
Walking			0.85						
<i>Other</i>				0.88			0.88		
Best specificity									
Drinking milk				0.98			0.98		
Grooming	0.99			0.99			0.99		
Lying		0.97						0.97	
Running	1.00	1.00		1.00	1.00		1.00	1.00	1.00
Walking	1.00			1.00			1.00		
<i>Other</i>			0.94						
Best precision									
Drinking milk							0.70		
Grooming				0.68			0.68		
Lying								0.95	
Running	0.95	0.95		0.95					
Walking				0.71					
<i>Other</i>					0.75				

3.2.3. Calculation time

Calculation time for features extraction, model training and testing are displayed in Table 8. While it took 3.8 minutes and 2.05 minutes to extract Catch22 and ROCKET features, respectively, more than 11.5 minutes were necessary for HC feature extraction. However, training of the RCV model was 203 and 64 times longer with ROCKET than with HC and Catch22, re-

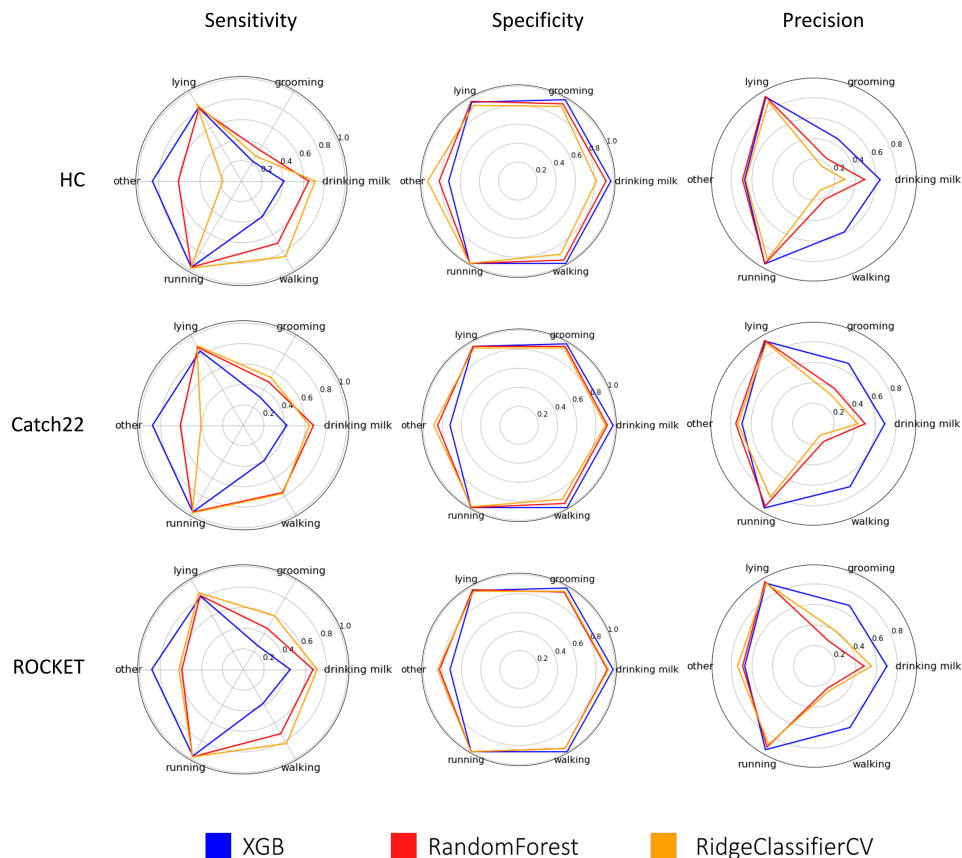


Figure 6: Sensitivity, specificity and precision performance per behaviour for every feature set-model combination.

spectively. Training XGB model took about five times longer with ROCKET features than with Catch22 and HC, respectively. Model testing took < 1 second regardless of the feature set-model combination. Based on the total calculation time, including features extraction, model training and testing, the fastest combination was ROCKET with RF (154.06 seconds), which yielded to the third-highest BA (0.71). The highest BA (0.77) obtained from ROCKET with RCV was x3.52 times slower than the fastest feature set-model combination but still took less than 10 minutes.



Figure 7: Confusion matrices for each feature-set model combination. Actual behaviours are in row; predicted behaviours are in column. Confusion between behaviours are expressed in percentage.

Table 8: Calculation time (seconds) for each feature set-model combination, including feature extraction, model training and testing.

		XGB	RandomForest	RidgeClassifierCV
Hand-Crafted	Feature Extraction		697.00	
	Training	491.86	42.75	2.06
	Testing	0.16	0.42	0.01
Catch22	Feature Extraction		228.00	
	Training	390.51	38.64	6.54
	Testing	0.17	0.58	0.01
ROCKET	Feature Extraction		123.00	
	Training	2232.42	30.47	420.00
	Testing	0.94	0.59	0.79

4. Discussion

4.1. Improving classification performance using ROCKET and Catch22

The main objective of that study was to evaluate the benefit of ROCKET and Catch22 features to classify pre-weaned calf behaviour from accelerometer data compared to a subset of features typically used in the field. Indeed, ROCKET and Catch22 are explicitly designed for time-series classification problems and exhibit impressive performance in related fields, justifying the evaluation of their performance in livestock ruminant behaviour classification. Therefore, we extracted ROCKET, Catch22 and HC features as the baseline, and tested the performance of the feature sets for the classification of 6 behaviours with 3 Machine Learning models, while considering calf inter-variability in performance evaluation. As expected, the best performance was achieved by ROCKET (0.70 ± 0.07 [mean \pm std]), closely followed by Catch22 (0.69 ± 0.05), both ranked ahead of HC (0.65 ± 0.03). Especially, the performances of the best combinations of models associated with Catch22 and ROCKET are significantly better than the best combination associated with HC. Indeed, Catch22 features achieved the highest BA of 0.73 using RF, which is a +5.8% improvement over the highest BA got with HC (0.69). Our study thereby confirms that Catch22 features are promising for classifying livestock ruminant behaviour from accelerometer time-series compared to HC features. Similarly, ROCKET features achieved the highest BA of 0.77 using RCV, which is a +11.59% improvement over the highest BA got with

HC (0.69). Considering that ROCKET can capture local and global patterns in the time-series and achieves better performance than cutting-edge classification techniques (Dempster et al., 2020), this finding is consistent with the literature and confirms that ROCKET features must be considered for classifying livestock ruminant behaviour from accelerometer data in further studies.

4.2. Relevance of feature sets to time-series

Running and lying performed well with ROCKET, Catch22 and HC features, combining a sensitivity, specificity and precision close to or higher than 0.80 regardless of the Machine Learning model. The singularity of those classes can explain this finding. As illustrated in Figure 8, lying behaviour leads to a flat signal without motion variation. In contrast, running behaviour leads to a high motion intensity signal with substantial variation on the three axes. The time-series signatures in the time-domain are highly specific, explaining why HC, Catch22 and ROCKET features are all able to discriminate lying and running from the *other* behaviours. However, although the specificity is usually higher than 0.80 whatever the behaviour and regardless of the set of features, there is a significant variation in the sensitivity and precision for drinking milk, grooming, walking and *other* depending on the feature set-model combination. HC and Catch22 features are more relevant for classifying walking, especially when combined with RCV and XGB, respectively. The walking time-series do not reveal any apparent shape in the time-domain (Figure 8). This finding is thus consistent with Lubba et al. (2019), who explain that time-series without time-aligned characteristic shapes are better suited for Catch 22 features-based representation. In contrast, ROCKET is more suitable for classifying grooming and drinking milk compared to Catch22 and HC, especially in association with RCV and XGB. As illustrated in Figure 8, grooming and drinking milk time-series have a subtle shape in the time-domain that are almost phase-aligned which is consistent with the temporal structure of these behaviours: Drinking milk is a repetitive sequence starting with milk suckling from the automatic feeder, following by swallowing, while grooming begins with a movement of the head to reach the area of the body to be licked, a sequence of licking movements and a return movement of the head. There is therefore a temporal structure to the movements of the jaw and head which must be reflected in the accelerometer time-series. Finding that ROCKET is more adapted than Catch22 or HC in that context is also coherent with Lubba

et al. (2019), who show that shape-based classifier, such as ROCKET and RCV, can accurately capture class differences in the time-series shape. Finally, the 3 features sets are relevant for discriminating the *other* behaviour. This class is the result of merging 19 behaviours, including oral manipulation of the pen, eating concentrates, sniffing, etc. This can be seen in Figure 8, where the time-series of the *other* class in the time-domain look highly different. As a result, some time-series may contain characteristic dynamic properties, while others exhibit subtle shape signatures, which explains why HC, Catch22 and ROCKET actively contribute to the classification of the *other* class. Therefore, our study confirms that different time-series classification problems require different time-series representations: Classes may be distinguished by multiple types of patterns and dynamical properties (Bagnall et al., 2017) and thus complementary features, such as Catch22 and ROCKET, are all necessary to achieve good performance in a multiple tasks classification problem.



Figure 8: Random 3 seconds time-windows selected from a set of random calves for each behaviour.

4.3. Impact of the Machine Learning model on the feature sets performance

Three Machine Learning models were tested in our study mainly to avoid any bias in evaluating feature sets performance. However, the results high-

lighted substantial variation within each set of features, suggesting that some Machine Learning models are more adapted to a given set. Based on the BA, the most performant models associated with ROCKET is RCV while the best performant model associated with Catch22 and HC is RF. Those results are consistent with the literature. Indeed, Dempster et al. (2020) state that even though ROCKET can be used with any classifier, it is very effective when used in conjunction with linear classifiers that can use a small amount of information from a large number of features. ROCKET and RCV are primarily effective due to their collaboration in managing high-dimensional feature spaces, fast cross-validation for regularization hyperparameters, and regularization mechanisms. This makes them a powerful combination for time-series classification tasks, especially in small datasets (Dempster et al., 2020; sktime Team, 2023; Middlehurst et al., 2023), as in ruminant behaviour classification, where the volume of data is often limited due to the manual annotation of behaviours (see section 4.5). Finding that RF is the best model associated with Catch22 and HC features is also consistent with the literature (Balli et al., 2019; Hu et al., 2020; Wang et al., 2020; Dickinson et al., 2021; Versluijs et al., 2023). RF is especially adequate for handling Catch22 and HC sets due to its ability to handle high-dimensional data and complex feature interactions. Its internal structure, consisting of multiple decision trees, enables robust and accurate modelling of diverse statistical features. Furthermore, even though ROCKET and RCV and Catch22 and RF are the best combinations for overall performance, other combinations look very promising to boost the performance metrics on some behaviours. Especially, Catch22 associated with XGB improves the precision of the behaviour drinking milk (0.69), grooming (0.68) and walking (0.71), which suffer from a large number of false positives with other combinations (see Appendix A). These results show that it is worth testing different models and choosing the one best suited to the features and the configuration of the study, as the best model is undoubtedly highly dependent on the behaviours to be classified, the volume of data, the class balance, etc.

4.4. *Limitations*

The best overall combination (ROCKET with RCV) reached a BA of 0.77, but improvements are still required to be used for the targeted applications. Also, the best overall combination does not produce the best performance per behaviour when all three metrics (sensitivity, specificity, precision) are considered. As explained above, playing with different fea-

ture set-model combinations is necessary to obtain the optimum for each behaviour. In addition, certain combinations can boost the sensitivity of the prediction, but as the expense of the precision. For example, ROCKET with RCV substantially boosted the sensitivity of grooming (0.60) but led to a high level of false positive (precision: 0.40). Similarly, ROCKET with XGB led to the highest precision for drinking milk (0.70) but also led to one of the lowest sensitivities (0.45). In addition, some behaviours lack sensitivity and/or precision, whatever the feature set-model combination (Table 7). In particular, drinking milk and grooming reached a maximum sensitivity of 0.60 and 0.71, and a precision of 0.68 and 0.70, respectively. Maximum precision of walking was also 0.71. These moderate performances are explained by the confusions between these behaviours (see Figure 7), probably due to the lack of characteristic dynamic properties and shapes in the accelerometer time-series collected from neck-collars.

These mitigated performances should be seen in light of the highly challenging scenario used in this study. First, six behaviours were kept in the classification to get a sufficient grain to assess calf welfare subsequently. However, the more classes are maintained in the classification, the poorer the performance (Riaboff et al., 2022). For example, Martiskainen et al. (2009) reached an average F1 score of 77.7% for 8 behaviours-classification using SVM while Benaissa et al. (2019) reached an accuracy of 92% for 3 behaviours-classification using the same model and accelerometer attachment system. Hosseininoorbin et al. (2021) observed a 6% improvement in the average F1 score when reducing the number of classes from 9 to 3. Specifically, the authors showed that performance decreases significantly when more than 5 behaviour classes are considered. This is also consistent with our study, where we observed a 7% improvement in the BA when decreasing the number of classes from 6 to 4 behaviours using ROCKET and RCV (*data not shown*). Furthermore, Hosseininoorbin et al. (2021) show that walking (F1-score: 32.19), drinking (F1-score: 8.85), grooming (F1-score: 15.60) and *other* (F1-score: 2.03) are extremely difficult to predict despite the robust architecture based on time-frequency data representation combined with Deep Neural Network. Those results are consistent with our study as walking, drinking milk and grooming are particularly difficult to classify, whatever the feature set-model combination. This can be explained by (i) a lack of data available to train and test the model, as these behaviours are rarely observed (Arablouei et al., 2023a) and (ii) a lack of characteristic dynamic properties or shapes that do not allow accurate classification. Secondly, the

performance in the study were evaluated according to a realistic but challenging scenario where the calves used for model testing have never been used for model training. Indeed, the decrement in model performance linked to high inter-animal variability has been widely reported in the literature (Riaboff et al., 2022). This is due to the differences in physical characteristics (e.g., muscles, tendons, joints) that lead to different expressions of the same behaviour (e.g., motion intensity, speed, posture (Barwick et al., 2020)). To support this, evaluation of ROCKET with RCV with a test set consisting of windows from a random split, irrespective of the calf from which they came, led to a 5% improvement in BA (*data not shown*). This trend is notably prominent for grooming and walking, where sensitivity increase (x1.24 and x1.34, respectively) with a random split (test_size=0.33). This is consistent with Hosseininoorbin et al. (2021), observing a 5% improvement in the F1-score when using random Stratified Cross Validation compared to Leave-One-Animal Cross Validation, with substantial improvement for the most poorly predicted behaviours.

4.5. Perspective

Regarding the limitations, it is worth noting that this study aimed to evaluate the potential of Catch22 and ROCKET features for calf behaviour classification from accelerometer data, considering their high performance for time-series classification in related fields. In that regard, no optimization has been implemented to boost performance. Assuming now that Catch22 and ROCKET are more performant than HC features, several techniques could be implemented to develop an accurate and robust classification system. First, dataset creation could be improved using the best feature set-model combination for each behaviour as a pre-trained model to identify the time-series where less prominent behaviours (drinking, grooming, walking) might occur. This could help to speed up the annotation process while improving the balance between classes, targeting directly the time-series and associated videos where the behaviours of interest may occur rather than annotating videos randomly. This should boost the classifier’s performance by enhancing data amount and quality, particularly for grooming and walking behaviours where there is few data available. Moreover, pre-processing could be further developed, investigating the benefit of low-pass filtering to remove signal noise and time-window overlap optimization (Riaboff et al., 2019). Regarding the modelling, feature selection could be applied to select the accelerometer time-series (raw axes, OBDA, pitch, roll, etc.) that are the most informative for

that classification task and the features associated with (Nanopoulos et al., 2001).

Furthermore, the study must be repeated with more data, different calves in test sets, etc., and if the limitations persist, a multi-vote classifier could be a promising way to improve the results. This multi-vote classifier can be based upon a system like the one suggested by Tahir et al. (2012) where they propose combining existing multi-label methods with ensemble techniques (heterogeneous ensemble). This can also help with sample imbalance and label correlation issues simultaneously. It also closely relates to the stacking idea, where the aggregation process is itself a meta-model (combining multiple machine learning models' predictions with another model enhances accuracy, leveraging the strengths of different base models for improved performance.) (Brownlee, 2021).

Another possible system would be based on binary classifiers. Smith et al. (2016) implemented a multi-class behaviour modelling based upon the "one-vs-all" framework: A set of binary classifiers was trained to discriminate one of the behaviour classes against all the remaining behaviours merged together. Confidence scores from the binary classifiers were combined to generate the estimated class. This approach led to a 5% performance improvement over standard multi-class time-series classifiers, creating diversity in the behaviour model. More recently, Arablouei et al. (2021) also obtained high performance with the "One versus Other" system, where 6 binary classifiers were trained to discriminate the class i versus class j . This approach could be implemented by selecting the best feature-model combinations for each binary classifier.

Finally, an assumption is that accelerometer data are not informative enough to discriminate behaviours with a fine grain regardless of the set of features. In that scenario, performance could be improved by combining accelerometer data with other sensors (Lee and Kwan, 2018; Cabezas et al., 2022; Arablouei et al., 2023b). Especially, combining accelerometer data with Computer Vision might help to improve the classification of behaviours which are dependent on the structure of the barn, such as drinking milk, eating concentrates, eating hay, etc. The complementary nature of the information provided by animal tracking (Vayssade et al., 2023) should boost substantially the detection of those behaviours while reducing the confusion with the other ones.

4.6. Practical impact of the study

Our study highlighted two sets of features, ROCKET and Catch22, designed explicitly in related communities for time-series classification problems. To the best of the authors' knowledge, Catch22 has never been applied in livestock ruminant classification from accelerometer data, while ROCKET has been implemented just once (Brouwers et al., 2023). Both ROCKET and Catch22 substantially improved the performance of calf behaviour classification compared to HC features, especially in association with RCV and RF, respectively. Catch22 and ROCKET are also fast to compute and easy to implement using ready-to-use libraries. Therefore, those features are extremely promising for the classification of livestock ruminant behaviour using accelerometer data and should be considered in future studies. In addition, to the authors' knowledge, this is one of the first studies to look at the classification of pre-weaned calf behaviour using accelerometer data, focusing on a wide range of behaviours and testing different methodologies (feature sets and models combinations). Moreover, while some improvement are necessary, performance could be easily optimized, resulting in a system suitable for experimental purposes. As mentioned in section 4.5, a heterogeneous ensemble approach or a binary classifier approach might be effective.

Such a behaviour classification system to monitor drinking milk, running, grooming, walking and lying activities could provide a global view of calf behaviour to support research in calf welfare. That could be used to measure the impact of routine practices (dehorning, transport, weaning from the dams, etc.) on animal behaviour to implement new practices that promote calf welfare in dairy farms. The system could also be applied to quantify the deviance in calf behaviour linked to sickness or distress. Such research may lead to the development of a decision tool to track each calf's behaviours individually and alert farmers to any anomalies detected. Further development would be, however, required to move from a system developed for experimental purposes to a practical system that farmers can use in commercial dairy farms (Arablouei et al., 2021).

5. Conclusion

This study aimed to evaluate the performance of ROCKET and Catch22 features for classifying 6 behaviours in pre-weaned calves, i.e., running, drinking milk, walking, lying, grooming and *other*. For that purpose, the performance of ROCKET and Catch22 features were compared to HC features

used in the field along with RF, XGB and RCV Machine Learning models. ROCKET and Catch22 substantially improved the BA compared to HC features. Especially the best feature set-model combination was ROCKET and RCV (BA: 0.77), followed by Catch22 and RF (BA: 0.73) and HC and RF (BA: 0.69). However, ROCKET and RCV do not yield the best sensitivity, specificity and precision for each behaviour. However, some of these variances could disappear with more data, which is for future evaluation. On the contrary, if the results suggest the same conclusions even after more data and optimisations, it can be hypothesised that several feature-set model combinations are necessary to get the optimal prediction of each behaviour separately. This supports the finding that different time-series classification problems require different time-series representations, and thus complementary features, such as Catch22 and ROCKET, are all necessary to perform well in classification. In that regard, a heterogeneous ensemble (multi-vote) system based on the most successful feature set-model combinations or a binary classifiers system would be highly valuable to boost the classification of each behaviour. Thus, a behaviour classification system produced as mentioned above could be used to evaluate the effect of practices applied routinely in pre-weaned dairy calves, such as dehorning, transport, weaning from the dams, etc., on calf behaviour to propose recommendations on practices promoting calf welfare in dairy farms.

Appendix A. Results in Detail

■ Best Result
 ■ Worst Result
 ■ Highest mean BA

Table A.9: Sensitivity values for each feature-model combination per behaviour.

Sensitivity							
		drinking_milk	grooming	lying	running	walking	other
Hand-Crafted	XGB	0.41	0.22	0.83	0.97	0.40	0.86
	Random Forest	0.65	0.35	0.82	0.97	0.70	0.61
	RidgeClassifierCV	0.71	0.28	0.87	0.98	0.85	0.18
	<i>Mean</i>	0.59	0.28	0.84	0.98	0.65	0.55
	<i>STD</i>	0.16	0.07	0.03	0.01	0.23	0.34
Catch22	XGB	0.42	0.32	0.84	0.98	0.40	0.88
	Random Forest	0.68	0.49	0.88	0.98	0.76	0.61
	RidgeClassifierCV	0.64	0.54	0.90	0.99	0.77	0.41
	<i>Mean</i>	0.58	0.45	0.87	0.98	0.64	0.63
	<i>STD</i>	0.14	0.11	0.03	0.00	0.21	0.24
Rocket	XGB	0.45	0.27	0.82	0.97	0.38	0.88
	Random Forest	0.67	0.46	0.83	0.98	0.72	0.59
	RidgeClassifierCV	0.71	0.60	0.86	0.98	0.83	0.62
	<i>Mean</i>	0.61	0.45	0.84	0.98	0.64	0.70
	<i>STD</i>	0.14	0.17	0.02	0.00	0.24	0.16

Table A.10: Specificity values for each feature-model combination per behaviour.

Specificity							
		drinking_milk	grooming	lying	running	walking	other
Hand-Crafted	XGB	0.97	0.99	0.96	1.00	1.00	0.72
	Random Forest	0.92	0.94	0.97	1.00	0.96	0.82
	RidgeClassifierCV	0.82	0.91	0.92	1.00	0.89	0.94
	<i>mean</i>	0.90	0.95	0.95	1.00	0.95	0.83
	<i>STD</i>	0.08	0.04	0.03	0.00	0.05	0.11
Catch22	XGB	0.98	0.99	0.96	1.00	1.00	0.72
	Random Forest	0.92	0.96	0.96	1.00	0.95	0.85
	RidgeClassifierCV	0.90	0.94	0.94	0.99	0.90	0.89
	<i>mean</i>	0.93	0.96	0.95	1.00	0.95	0.82
	<i>STD</i>	0.04	0.03	0.01	0.00	0.05	0.09
Rocket	XGB	0.98	0.99	0.96	1.00	1.00	0.72
	Random Forest	0.92	0.94	0.97	1.00	0.96	0.83
	RidgeClassifierCV	0.93	0.95	0.95	1.00	0.96	0.85
	<i>mean</i>	0.94	0.96	0.96	1.00	0.97	0.80
	<i>STD</i>	0.03	0.03	0.01	0.00	0.02	0.07

Table A.11: Precision values for each feature-model combination per behaviour.

Precision							
		drinking_milk	grooming	lying	running	walking	other
Hand-Crafted	XGB	0.64	0.46	0.93	0.95	0.59	0.67
	Random Forest	0.49	0.24	0.94	0.95	0.22	0.69
	RidgeClassifierCV	0.30	0.15	0.88	0.90	0.12	0.66
	<i>mean</i>	0.48	0.28	0.92	0.93	0.31	0.67
	<i>STD</i>	0.17	0.16	0.03	0.03	0.25	0.02
Catch22	XGB	0.69	0.68	0.93	0.95	0.71	0.69
	Random Forest	0.50	0.40	0.94	0.93	0.20	0.75
	RidgeClassifierCV	0.43	0.33	0.91	0.83	0.13	0.73
	<i>mean</i>	0.54	0.47	0.92	0.90	0.35	0.72
	<i>STD</i>	0.14	0.19	0.02	0.06	0.32	0.03
Rocket	XGB	0.70	0.68	0.93	0.94	0.69	0.67
	Random Forest	0.48	0.28	0.95	0.91	0.25	0.69
	RidgeClassifierCV	0.55	0.40	0.93	0.88	0.28	0.74
	<i>mean</i>	0.58	0.45	0.94	0.91	0.41	0.70
	<i>STD</i>	0.11	0.21	0.01	0.03	0.25	0.03

References

- AlZubi, H.S., Gerrard-Longworth, S., Al-Nuaimy, W., Goulermas, Y., Preece, S., 2014. Human activity classification using a single accelerometer, in: 2014 14th UK Workshop on Computational Intelligence (UKCI), IEEE. pp. 1–6.
- Arablouei, R., Currie, L., Kusy, B., Ingham, A., Greenwood, P.L., Bishop-Hurley, G., 2021. In-situ classification of cattle behavior using accelerometry data. *Computers and Electronics in Agriculture* 183, 106045.
- Arablouei, R., Wang, L., Currie, L., Yates, J., Alvarenga, F.A., Bishop-Hurley, G.J., 2023a. Animal behavior classification via deep learning on embedded systems. *Computers and Electronics in Agriculture* 207, 107707.
- Arablouei, R., Wang, Z., Bishop-Hurley, G.J., Liu, J., 2023b. Multimodal sensor data fusion for in-situ classification of animal behavior using accelerometry and gns data. *Smart Agricultural Technology* 4, 100163.

- Aziz, W., Hussain, L., Khan, I.R., Alowibdi, J.S., Alkinani, M.H., 2021. Machine learning based classification of normal, slow and fast walking by extracting multimodal features from stride interval time series. *Math. Biosci. Eng* 18, 495–517.
- Bagnall, A., Lines, J., Bostrom, A., Large, J., Keogh, E., 2017. The great time series classification bake off: a review and experimental evaluation of recent algorithmic advances. *Data mining and knowledge discovery* 31, 606–660.
- Balasso, P., Taccioli, C., Serva, L., Magrin, L., Andrighetto, I., Marchesini, G., 2023. Uncovering patterns in dairy cow behaviour: A deep learning approach with tri-axial accelerometer data. *Animals* 13, 1886.
- Balli, S., Sağbaşı, E.A., Peker, M., 2019. Human activity recognition from smart watch sensor data using a hybrid of principal component analysis and random forest algorithm. *Measurement and Control* 52, 37–45.
- Barry, J., Kennedy, E., Sayers, R., de Boer, I.J., Bokkers, E., 2019. Development of a welfare assessment protocol for dairy calves from birth through to weaning. *Animal Welfare* 28, 331–344.
- Barwick, J., Lamb, D.W., Dobos, R., Welch, M., Schneider, D., Trotter, M., 2020. Identifying sheep activity from tri-axial acceleration signals using a moving window classification model. *Remote Sensing* 12, 646.
- Barwick, J., Lamb, D.W., Dobos, R., Welch, M., Trotter, M., 2018. Categorising sheep activity using a tri-axial accelerometer. *Computers and Electronics in Agriculture* 145, 289–297.
- Benaissa, S., Tuytens, F.A., Plets, D., De Pessemier, T., Trogh, J., Tanghe, E., Martens, L., Vandaele, L., Van Nuffel, A., Joseph, W., et al., 2019. On the use of on-cow accelerometers for the classification of behaviours in dairy barns. *Research in veterinary science* 125, 425–433.
- Bersch, S.D., Azzi, D., Khusainov, R., Achumba, I.E., Ries, J., 2014. Sensor data acquisition and processing parameters for human activity classification. *Sensors* 14, 4239–4270.
- Breiman, L., 2001. Random forests. *Machine learning* 45, 5–32.

- Brouwers, S.P., Simmler, M., Savary, P., Scriba, M.F., 2023. Towards a novel method for detecting atypical lying down and standing up behaviors in dairy cows using accelerometers and machine learning. *Smart Agricultural Technology* 4, 100199.
- Brownlee, J., 2021. Ensemble learning algorithms with Python: Make better predictions with bagging, boosting, and stacking. *Machine Learning Mastery*.
- Cabezas, J., Yubero, R., Visitación, B., Navarro-García, J., Algar, M.J., Cano, E.L., Ortega, F., 2022. Analysis of accelerometer and gps data for cattle behaviour identification and anomalous events detection. *Entropy* 24, 336.
- Cardoso, C.S., Hötzel, M.J., Weary, D.M., Robbins, J.A., von Keyserlingk, M.A., 2016. Imagining the ideal dairy farm. *Journal of Dairy Science* 99, 1663–1671.
- Carslake, C., Vázquez-Diosdado, J.A., Kaler, J., 2020. Machine learning algorithms to classify and quantify multiple behaviours in dairy calves using a sensor: Moving beyond classification in precision livestock. *Sensors* 21, 88.
- Chakravarty, P., Cozzi, G., Ozgul, A., Aminian, K., 2019. A novel biomechanical approach for animal behaviour recognition using accelerometers. *Methods in Ecology and Evolution* 10, 802–814.
- Chang, A.Z., Fogarty, E.S., Moraes, L.E., García-Guerra, A., Swain, D.L., Trotter, M.G., 2022. Detection of rumination in cattle using an accelerometer ear-tag: A comparison of analytical methods and individual animal and generic models. *Computers and Electronics in Agriculture* 192, 106595.
- Chen, T., Guestrin, C., 2016. Xgboost: A scalable tree boosting system, in: *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining*, pp. 785–794.
- Conneely, M., Berry, D., Murphy, J., Lorenz, I., Doherty, M., Kennedy, E., 2014. Effects of milk feeding volume and frequency on body weight and health of dairy heifer calves. *Livestock Science* 161, 90–94.

- Dempster, A., Petitjean, F., Webb, G.I., 2020. Rocket: exceptionally fast and accurate time series classification using random convolutional kernels. *Data Mining and Knowledge Discovery* 34, 1454–1495.
- Dempster, A., Schmidt, D.F., Webb, G.I., 2021. Minirocket: A very fast (almost) deterministic transform for time series classification, in: *Proceedings of the 27th ACM SIGKDD conference on knowledge discovery & data mining*, pp. 248–257.
- Dickinson, E.R., Twining, J.P., Wilson, R., Stephens, P.A., Westander, J., Marks, N., Scantlebury, D.M., 2021. Limitations of using surrogates for behaviour classification of accelerometer data: refining methods using random forest models in caprids. *Movement Ecology* 9, 28.
- Dissanayake, O., McPherson, S., Kennedy, E., Sugrue, K., Conneely, M., Shaloo, L., Cunningham, P., Riaboff, L., 2022. Identification of the best accelerometer features and time-scale to detect disturbances in calves, in: *International Workshop on Advanced Analytics and Learning on Temporal Data*, Springer. pp. 167–180.
- Drover, D., Howcroft, J., Kofman, J., Lemaire, E.D., 2017. Faller classification in older adults using wearable sensors based on turn and straight-walking accelerometer-based features. *Sensors* 17, 1321.
- Duthie, C.A., Bowen, J., Bell, D., Miller, G., Mason, C., Haskell, M., 2021. Feeding behaviour and activity as early indicators of disease in pre-weaned dairy calves. *Animal* 15, 100150.
- Dutta, R., Smith, D., Rawnsley, R., Bishop-Hurley, G., Hills, J., Timms, G., Henry, D., 2015. Dynamic cattle behavioural classification using supervised ensemble classifiers. *Computers and electronics in agriculture* 111, 18–28.
- Enríquez, D., Ungerfeld, R., Quintans, G., Guidoni, A., Hötzel, M., 2010. The effects of alternative weaning methods on behaviour in beef calves. *Livestock Science* 128, 20–27.
- Fan, S., Jia, Y., Jia, C., 2019. A feature selection and classification method for activity recognition based on an inertial sensing unit. *Information* 10, 290.

- Fida, B., Bernabucci, I., Bibbo, D., Conforto, S., Schmid, M., 2015. Pre-processing effect on the accuracy of event-based activity segmentation and classification through inertial sensors. *Sensors* 15, 23095–23109.
- Figo, D., Diniz, P.C., Ferreira, D.R., Cardoso, J.M., 2010. Preprocessing techniques for context recognition from accelerometer data. *Personal and Ubiquitous Computing* 14, 645–662.
- Fogarty, E.S., Swain, D.L., Cronin, G.M., Moraes, L.E., Trotter, M., 2020. Behaviour classification of extensively grazed sheep using machine learning. *Computers and Electronics in Agriculture* 169, 105175.
- Friard, O., Gamba, M., 2016. Boris: a free, versatile open-source event-logging software for video/audio coding and live observations. *Methods in ecology and evolution* 7, 1325–1330.
- Friedman, J.H., 2001. Greedy function approximation: a gradient boosting machine. *Annals of statistics* , 1189–1232.
- González, L.A., Bishop-Hurley, G., Handcock, R.N., Crossman, C., 2015. Behavioral classification of data from collars containing motion sensors in grazing cattle. *Computers and electronics in agriculture* 110, 91–102.
- van Hees, V.T., Gorzelniak, L., Dean León, E.C., et al., 2013. Separating movement and gravity components in an acceleration signal and implications for the assessment of human daily physical activity. *PLoS One* 8, e61691. URL: <https://doi.org/10.1371/journal.pone.0061691>, doi:10.1371/journal.pone.0061691. published 2013 Apr 23.
- Hoerl, A.E., Kennard, R.W., 1970. Ridge regression: Biased estimation for nonorthogonal problems. *Technometrics* 12, 55–67.
- Hosseininoorbin, S., Layeghy, S., Kusy, B., Jurdak, R., Bishop-Hurley, G.J., Greenwood, P.L., Portmann, M., 2021. Deep learning-based cattle behaviour classification using joint time-frequency data representation. *Computers and Electronics in Agriculture* 187, 106241.
- Hounslow, J., Brewster, L., Lear, K., Guttridge, T., Daly, R., Whitney, N., Gleiss, A., 2019. Assessing the effects of sampling frequency on behavioural classification of accelerometer data. *Journal of experimental marine biology and ecology* 512, 22–30.

- Hu, S., Ingham, A., Schmoelzl, S., McNally, J., Little, B., Smith, D., Bishop-Hurley, G., Wang, Y.G., Li, Y., 2020. Inclusion of features derived from a mixture of time window sizes improved classification accuracy of machine learning algorithms for sheep grazing behaviours. *Computers and Electronics in Agriculture* 179, 105857.
- Iqbal, M.W., Draganova, I., Morel, P.C., Morris, S.T., 2021. Validation of an accelerometer sensor-based collar for monitoring grazing and rumination behaviours in grazing dairy cows. *Animals* 11, 2724.
- Jiang, B., Tang, W., Cui, L., Deng, X., 2023. Precision livestock farming research: A global scientometric review. *Animals* 13, 2096.
- Kleanthous, N., Hussain, A.J., Khan, W., Sneddon, J., Al-Shamma'a, A., Liatsis, P., 2022. A survey of machine learning approaches in animal behaviour. *Neurocomputing* 491, 442–463.
- Koknaroglu, H., Akunal, T., 2013. Animal welfare: An animal science approach. *Meat Science* 95, 821–827.
- Lee, K., Kwan, M.P., 2018. Automatic physical activity and in-vehicle status classification based on gps and accelerometer data: A hierarchical classification approach using machine learning techniques. *Transactions in GIS* 22, 1522–1549.
- Löning, M., Bagnall, A., Ganesh, S., Kazakov, V., Lines, J., Király, F.J., 2019. sktime: A unified interface for machine learning with time series. *arXiv preprint arXiv:1909.07872* .
- Lubba, C.H., Sethi, S.S., Knaute, P., Schultz, S.R., Fulcher, B.D., Jones, N.S., 2019. catch22: Canonical time-series characteristics: Selected through highly comparative time-series analysis. *Data Mining and Knowledge Discovery* 33, 1821–1852.
- Lush, L., Wilson, R.P., Holton, M.D., Hopkins, P., Marsden, K.A., Chadwick, D.R., King, A.J., 2018. Classification of sheep urination events using accelerometers to aid improved measurements of livestock contributions to nitrous oxide emissions. *Computers and electronics in agriculture* 150, 170–177.

- Mahendran, S.A., Wathes, D.C., Blackie, N., Booth, R.E., 2023. Calf health, feeding and social behaviours within groups fed on automatic milk feeders. *Ruminants* 3, 266–285.
- Martiskainen, P., Järvinen, M., Skön, J.P., Tiirikainen, J., Kolehmainen, M., Mononen, J., 2009. Cow behaviour pattern recognition using a three-dimensional accelerometer and support vector machines. *Applied animal behaviour science* 119, 32–38.
- McPherson, S., Riaboff, L., Dissanayake, O., Sinnott, A., Cunningham, P., Kennedy, E., 2022. Effect of separation at weaning on the activity of cows and calves reared in a cow-calf contact system measured with accelerometer sensors, in: 10th European Conference on Precision Livestock Farming, Vienna.
- Middlehurst, M., Schäfer, P., Bagnall, A., 2023. Bake off redux: a review and experimental evaluation of recent time series classification algorithms. *arXiv preprint arXiv:2304.13029* .
- Mietus, J., Peng, C., Henry, I., Goldsmith, R., Goldberger, A., 2002. The pnnx files: re-examining a widely used heart rate variability measure. *Heart* 88, 378.
- Moreau, M., Siebert, S., Buerkert, A., Schlecht, E., 2009. Use of a tri-axial accelerometer for automated recording and classification of goats' grazing behaviour. *Applied animal behaviour science* 119, 158–170.
- Morisse, J., Cotte, J., Huonnic, D., 1995. Effect of dehorning on behaviour and plasma cortisol responses in young calves. *Applied Animal Behaviour Science* 43, 239–247.
- Nanopoulos, A., Alcock, R., Manolopoulos, Y., 2001. Feature-based classification of time-series data. *International Journal of Computer Research* 10, 49–61.
- Nikkhah, A., Alimirzaei, M., 2023. Understanding calf behavioral responses to environmental changes and challenges: An applied update. *Farm Animal Health and Nutrition* 2, 71–77.
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J.,

- Passos, A., Cournapeau, D., Brucher, M., Perrot, M., Duchesnay, E., 2011. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research* 12, 2825–2830.
- Penning, P., 1983. A technique to record automatically some aspects of grazing and ruminating behaviour in sheep. *Grass and Forage Science* 38, 89–96.
- Preece, S.J., Goulermas, J.Y., Kenney, L.P., Howard, D., 2008. A comparison of feature extraction methods for the classification of dynamic activities from accelerometer data. *IEEE Transactions on Biomedical Engineering* 56, 871–879.
- Price, E., Langford, J., Fawcett, T.W., Wilson, A.J., Croft, D.P., 2022. Classifying the posture and activity of ewes and lambs using accelerometers and machine learning on a commercial flock. *Applied Animal Behaviour Science* 251, 105630.
- Rahman, A., Smith, D., Little, B., Ingham, A., Greenwood, P., Bishop-Hurley, G., 2018. Cattle behaviour classification from collar, halter, and ear tag sensors. *Information processing in agriculture* 5, 124–133.
- Riaboff, L., Aubin, S., Bédère, N., Couvreur, S., Madouasse, A., Goumand, E., Chauvin, A., Plantier, G., 2019. Evaluation of pre-processing methods for the prediction of cattle behaviour from accelerometer data. *Computers and electronics in agriculture* 165, 104961.
- Riaboff, L., Poggi, S., Madouasse, A., Couvreur, S., Aubin, S., Bédère, N., Goumand, E., Chauvin, A., Plantier, G., 2020. Development of a methodological framework for a robust prediction of the main behaviours of dairy cows using a combination of machine learning algorithms on accelerometer data. *Computers and Electronics in Agriculture* 169, 105179.
- Riaboff, L., Shalloo, L., Smeaton, A.F., Couvreur, S., Madouasse, A., Keane, M.T., 2022. Predicting livestock behaviour using accelerometers: A systematic review of processing techniques for ruminant behaviour prediction from raw accelerometer data. *Computers and Electronics in Agriculture* 192, 106610.

- Rodriguez-Baena, D.S., Gomez-Vela, F.A., García-Torres, M., Divina, F., Barranco, C.D., Daz-Diaz, N., Jimenez, M., Montalvo, G., 2020. Identifying livestock behavior patterns based on accelerometer dataset. *Journal of Computational Science* 41, 101076.
- Rushen, J., Chapinal, N., de Passile, A., 2012. Automated monitoring of behavioural-based animal welfare indicators. *Animal Welfare* 21, 339–350.
- da Silva, I.C., van Hees, V.T., Ramires, V.V., Knuth, A.G., Bielemann, R.M., Ekelund, U., Brage, S., Hallal, P.C., 2014. Physical activity levels in three brazilian birth cohorts as assessed with raw triaxial wrist accelerometry. *International journal of epidemiology* 43, 1959–1968.
- Smith, D., Rahman, A., Bishop-Hurley, G.J., Hills, J., Shahriar, S., Henry, D., Rawnsley, R., 2016. Behavior classification of cows fitted with motion collars: Decomposing multi-class classification into a set of binary problems. *Computers and Electronics in Agriculture* 131, 40–50.
- Tahir, M.A., Kittler, J., Bouridane, A., 2012. Multilabel classification using heterogeneous ensemble of multi-label classifiers. *Pattern Recognition Letters* 33, 513–523.
- Tatler, J., Cassey, P., Prowse, T.A., 2018. High accuracy at low frequency: detailed behavioural classification from accelerometer data. *Journal of Experimental Biology* 221, jeb184085.
- sktime Team, 2023. Demo of rocket transform — sktime documentation. <https://www.sktime.org/en/latest/examples/rocket.html>.
- Vayssade, J.A., Godard, X., Bonneau, M., 2023. Wizard: Unsupervised goats tracking algorithm. *Computers and Electronics in Agriculture* 209, 107831.
- Vázquez Diosdado, J.A., Barker, Z.E., Hodges, H.R., Amory, J.R., Croft, D.P., Bell, N.J., Codling, E.A., 2015. Classification of behaviour in housed dairy cows using an accelerometer-based activity monitoring system. *Animal Biotelemetry* 3, 1–14.
- Versluijs, E., Niccolai, L.J., Spedener, M., Zimmermann, B., Hesse, A., Tofastrud, M., Devineau, O., Evans, A.L., 2023. Classification of behaviors of free-ranging cattle using accelerometry signatures collected by virtual fence collars. *Frontiers in Animal Science* 4, 1083272.

- Walker, J.S., Jones, M.W., Laramée, R.S., Holton, M.D., Shepard, E.L., Williams, H.J., Scantlebury, D.M., Marks, N.J., Magowan, E.A., Maguire, I.E., et al., 2015. Prying into the intimate secrets of animal lives; software beyond hardware for comprehensive annotation in ‘daily diary’ tags. *Movement ecology* 3, 1–16.
- Wang, A., Chen, H., Zheng, C., Zhao, L., Liu, J., Wang, L., 2020. Evaluation of random forest for complex human activity recognition using wearable sensors, in: 2020 International Conference on Networking and Network Applications (NaNA), IEEE. pp. 310–315.
- Wang, X., Wirth, A., Wang, L., 2007. Structure-based statistical features and multivariate time series clustering, in: Seventh IEEE international conference on data mining (ICDM 2007), IEEE. pp. 351–360.
- Wilson, R.P., Shepard, E., Liebsch, N., 2008. Prying into the intimate details of animal lives: use of a daily diary on animals. *Endangered species research* 4, 123–137.
- Yu, H., Muijres, F.T., te Lindert, J.S., Hedenström, A., Henningsson, P., 2023. Accelerometer sampling requirements for animal behaviour classification and estimation of energy expenditure. *Animal Biotelemetry* 11, 28.
- Zdravevski, E., Risteska Stojkoska, B., Standl, M., Schulz, H., 2017. Automatic machine-learning based identification of jogging periods from accelerometer measurements of adolescents under field conditions. *PloS one* 12, e0184216.