Paper:

Wilson, R., Holton, M., di Virgilio, A., Williams, H., Shepard, E., Lambertucci, S., Quintana, F., Sala, J., Balaji, B., et. al. (2018). Give the machine a hand: A Boolean time-based decision-tree template for rapidly finding animal behaviours in multisensor data. *Methods in Ecology and Evolution*

http://dx.doi.org/10.1111/2041-210X.13069

Give the machine a hand: A Boolean time-based decision-tree template for rapidly finding animal behaviours in multi-sensor data

Rory P. Wilson[1(ID)], Mark D. Holton[2], Agustina di Virgilio[3,4], Hannah Williams[1], Emily L. C. Shepard[1], Sergio Lambertucci[3], Flavio Quintana[5], Juan E. Sala[5], Bharathan Balaji[6], Eun Sun Lee[6], Mani Srivastava[6], D. Michael Scantlebury[7], Carlos M. Duarte[8].

[1]*Department of Biosciences, College of Science, Swansea University, Swansea SA2 8PP, UK*
[2]*Department of Computing Science, College of Science, Swansea University, Swansea SA2 8PP, UK*
[3]*Grupo de Biología de la Conservación, Laboratorio Ecotono, INIBIOMA (CONICET-Universidad Nacional del Comahue), Bariloche, Argentina.*
[4]*Grupo de Ecología Cuantitativa, INIBIOMA (CONICET-Universidad Nacional del Comahue), Bariloche, Argentina*
[5]*Instituto de Biologia de Organismos Marinos IBIOMAR-CONICET (9120) Puerto Madryn, Chubut, Argentina*
[6]*Department of Electrical and Computer Engineering, University of California, Los Angeles, Los Angeles, California 90095, USA*
[7]*School of Biological Sciences, Institute for Global Food Security, Queen's University Belfast, Belfast BT9 7BL, UK*
[8]*Red Sea Research Centre, King Abdullah University of Science and Technology, Thuwal 23955, Saudi Arabia*

Abstract

1. The development of multi-sensor animal-attached tags, recording data at high frequencies, has enormous potential in allowing us to define animal behaviour.
2. The high volumes of data, are pushing us towards machine-learning as a powerful option for distilling out behaviours. However, with increasing parallel lines of data, systems

become more likely to become processor limited and thereby take appreciable amounts of time to resolve behaviours.

3. We suggest a Boolean approach whereby critical changes in recorded parameters are used as sequential templates with defined flexibility (in both time and degree) to determine individual behavioural elements within a behavioural sequence that, together, makes up a single, defined behaviour.
4. We tested this approach, and compared it to a suite of other behavioural identification methods, on a number of behaviours from tag-equipped animals; sheep grazing, penguins walking, cheetah stalking prey and condors thermalling.
5. Overall behaviour recognition using our new approach was better than most other methods due to; (i) its ability to deal with behavioural variation and (ii) the speed with which the task was completed because extraneous data are avoided in the process.
6. We suggest that this approach is a promising way forward in an increasingly data-rich environment and that workers sharing algorithms can provide a powerful library for the benefit of all involved in such work.

1 | INTRODUCTION

Animal behaviour has been variously defined, but generally can be defined as 'the way in which an animal works, functions or responds to a particular situation' (Tinbergen 1960) with consequences for lifetime reproductive success (Birkhead, Atkin & Møller 1987; Drews 1993; Krebs, Davies & Parr 1993; Krebs & Davies 2009). As such, our ability to determine animal behaviours precisely is critically important for proper understanding of animal ecology and ecosystem functioning (Krebs, Davies & Parr 1993). Indeed, it is this that explains the large variety of methodologies developed to quantify behaviour (e.g. Tinbergen 1960; Altmann 1974; Lucas & Baras 2000; Miller & Gerlai 2007; Chastin & Granat 2010). A particularly rapidly developing field in this regard is 'biologging' – the deployment of autonomous tags on animals to record data (Hooker *et al.* 2007). Specifically, the extraordinary development of electronic technology over the last 3 decades has led the progression of sophisticated miniature sensors coupled with low power consumption and rapidly expanding memory capacity (Ropert-Coudert & Wilson 2005) so that studies using multi-sensor technology in tags on animals are now common (Brown *et al.* 2013). This has led from the simple animal-attached tags of the 1990s recording data once every few seconds (Wilson *et al.* 1994), to systems today that may record multiple channels at thousands of Hertz (Johnson & Tyack 2003). Of particular note for defining behaviours is the role played by accelerometers, gyroscopes and magnetometers, which can resolve both animal attitude in the 3 spatial axes (Yoda *et al.* 1999; Williams *et al.* 2017) and movement (Fourati *et al.* 2011; Noda *et al.* 2014). These are primary elements used in classifying behaviours (Tinbergen 1960), and so have great potential in studies of wild animals.

However, the ease with which we can now record the physical manifestation of behaviour, *via* metrics such as pitch, roll and 'dynamism' in the acceleration signature (Laich *et al.* 2008), is tempered by the difficulties of dealing with the complexity and volume of such data. Thus, computational solutions for processing the signals are inevitable and, accordingly, there is a rich and varied literature dealing with this (e.g. Sakamoto *et al.* 2009; Nathan *et al.* 2012; Resheff *et al.* 2014). This includes support vector machines (Tachibana, Oosugi & Okanoya 2014), regression trees (de Weerd *et al.* 2015), random forests (Bidder *et al.* 2014), neural networks (Samarasinghe 2016), linear discriminant analysis (Anderberg 2014) and template-matching (Walker *et al.* 2015b). Each method has advantages and disadvantages (Resheff *et al.* 2014) but prime negative issues revolve around subjectivity, whether the data are parametric, the extent of over-fitting, and the computational time involved in the process (Nathan *et al.* 2012). In addition, a particular weakness of many systems is that they fail to recognise the temporal sequencing of the movements that define the fundamental unit of that behaviour and the variability within them, and thereby preclude an important discriminator. For example, walking may be defined by a cluster of acceleration metrics (Bidder *et al.* 2014) but the fundamental unit of walking is the single step (Moe-Nilssen & Helbostad 2004) and this has well-defined properties over time (Sabatini *et al.* 2005) that could, for example, be used in any decision tree-based approach.

In this paper, we present an approach for identifying behaviours from data derived from animal-attached tags that recognises (i) the lowest common denominator (LoCoD) defining any particular behaviour (i.e. a single step is the lowest common denominator within walking) and (ii) that this lowest common denominator can be usefully broken down into base elements (BEs) (such as an increase, followed by a drop, in dorso-ventral acceleration for walking (Rong *et al.* 2007)), all of which have to follow each other in a defined sequence for the LoCoD to be apparent. Finally, (iii), the timing of BEs within a sequence is often constrained. Thus, this process provides a recognizable key for LoCoDs of behaviours based on measurements, sequences and timings of BEs. We appreciate that much of the essence of this is inherent in some template-matching approaches (Walker *et al.* 2015a) but combine this with both temporal flexibility across all BEs, together with an ability to switch between and incorporate defined, often derived, metrics that provide critical information for a powerful match. We demonstrate the utility of this approach by using it to search for behaviours that have LoCoD periods ranging between fractions of a second and several minutes using data derived from animal-attached tags and compare it briefly to other computational methods.

2 | MATERIALS AND METHODS

For this approach, we consider primary data derived from orthogonal, tri-axial accelerometers as well as, where helpful, information from pressure- and magnetic sensors, in addition to calculated variables obtained from acceleration data, such as Vectorial Dynamic Body Acceleration (VeDBA) (Qasem *et al.* 2012).

2.1 | The LoCoD Method

The LoCoD method involves initial consideration of the data visually by the user, who should examine the details of the movement that makes up the behaviour and reflect how this movement is expected to affect the sensors. In this, the user should identify the patterns that make up the BEs of the LoCoD and whether they can be made more distinctive by selective smoothing, as is done in many behaviour-identifying protocols anyway (Nathan *et al.* 2012). In addition, it is recommended that differentials be derived for any signals of interest, since these often act as excellent thresholds in derivation of the BEs (Fig. 1). Differentials are particularly important since postural data derived from acceleration (Shepard *et al.* 2008) are dependent, in part, on the angle of the terrain beneath the study animal (cf. the difference in sway axis during the stationary periods at the beginning and end of the walking period in Fig. 1), as well as the tag placement. Thus, working with differentials essentially standardises the signal output.

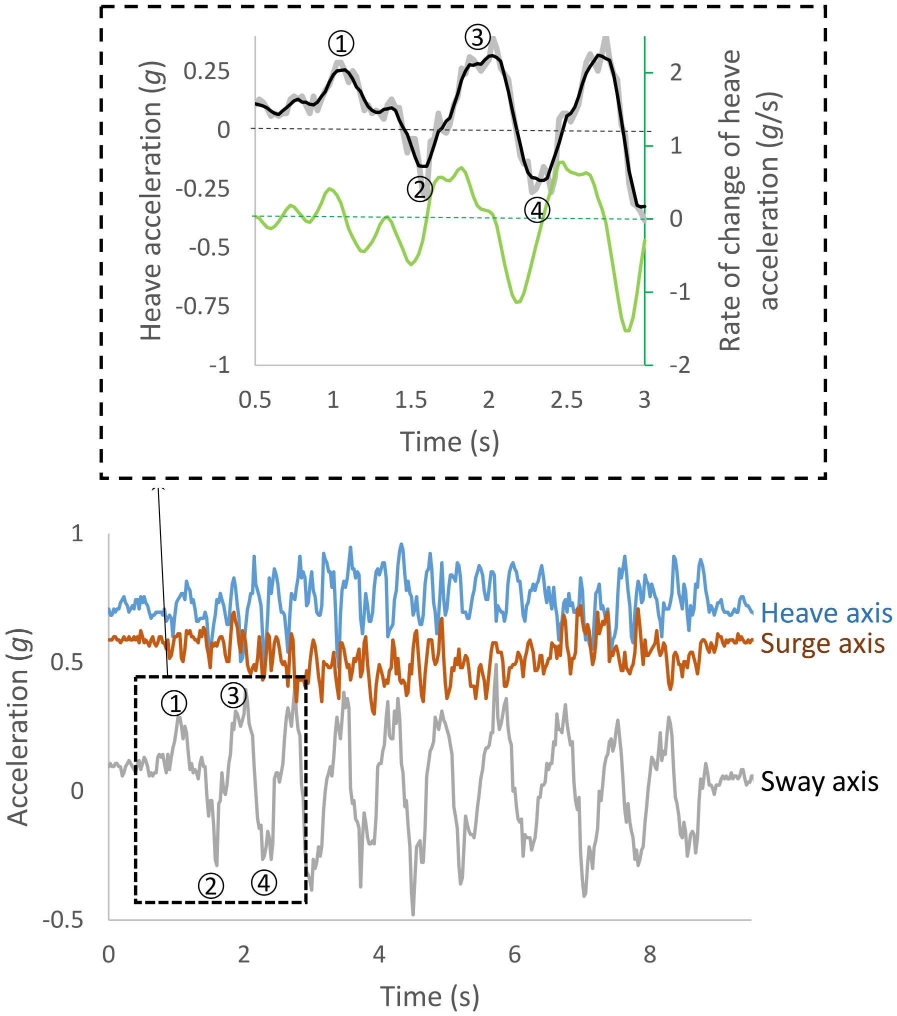

Fig. 1 – Twenty steps (the first 4 numbered) taken by a Magellanic penguin *Spheniscus magellanicus* during walking on the beach, manifest by tri-axial acceleration data at 40 Hz. The bird starts and ends stationary, but begins to walk, with 2 small steps before rapidly changing to steps with clear waveforms, particularly in the sway (lateral) axis (grey line). Within the LoCoD framework, the user is expected to identify the most useful primary data streams for the process. These may be expanded by deriving secondary data streams, such as smoothed values, to enhance BE identification. The inset shows the first 5 steps (grey line) smoothed over 0.125 s (black line) in the dominant waveform (the sway axis) and the rate of change of the smoothed data (green line).

Following decisions on which channels are to be used for identification of the behaviour, the conditions describing each BE are set up in ordered sequence to describe the LoCoD. Each summary condition for the BE follows a Boolean approach. For example, summary condition 1 that defines BE_1 of the LoCoD for a penguin walking (Fig. 2) may be asked to recognise the moment when the differential of the smoothed sway acceleration exceeds 0.25 *g*/s;

$$BE_1 \text{ - RECOGNISE WHEN; } dAh_s/dt > 0.25\ g/s \qquad (1)$$

where Ah_s is the smoothed heave acceleration following;

$$Ahs = \frac{1}{n} \sum_{i=0}^{n-1} Ah - i \qquad (2)$$

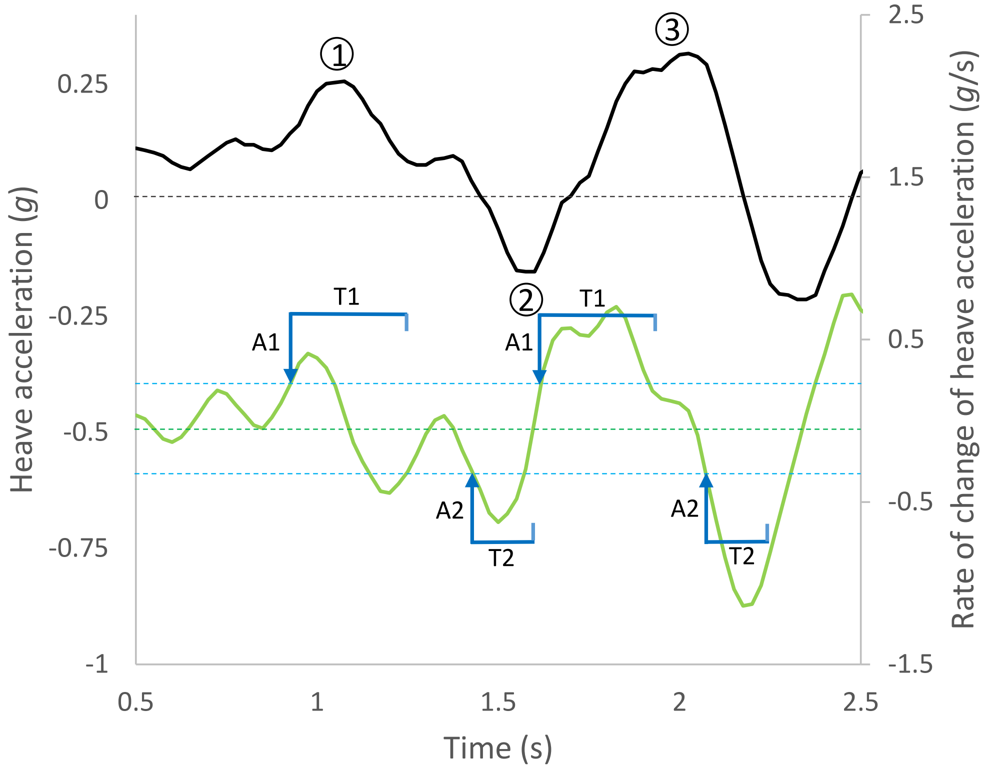

Fig. 2 – The first 3 steps (numbered) of the walking period shown in Fig. 1 for the smoothed heave axis (black line) and the rate of change of the heave axis (green line). The LoCoD method first identifies a feature, or combination of features, that signify the initiation of the first BE of the behaviour (here a differential threshold of >0.25 *g*/s) (marked A1). There is then a defined 'dead' time (T1), over which the program skips before looking for the second BE defining the behaviour (here a differential threshold of < -0.25 *g*/s) (A2) with its 'dead' time (T2). If these two conditions are met (as in this case) the LoCoD is made of 2 BEs and describes the conditions for one left stride followed by one right stride. The process could, however, be used for strides from one leg only, for example, whereupon either just A1 and T1 or A2 and T2 would be used for left and right strides, respectively.

In addition, the process should recognise multiple, cross-channel sub-conditions (for positives and negatives). Thus, equation (1) might be made of 3 sub-conditions;

BE_1 - RECOGNISE WHEN; $dAh_s/dt > 0.25$ *g*/s

AND; $dAs/dt > 0.05\ g/s$

AND NOT; $D > 0$ m (3)

where As is the surge acceleration and D is the depth.

Importantly, each sub-condition or condition can employ a time base with three elements within it that can be specified. These are;

1. *Presence* - that the sub-condition or condition is maintained over a specified time for the statement to be TRUE
2. *Range* - that, following identification of a true sub-condition or condition, the program can skip a defined number of data points before looking for the next BE. This is important because it can stop the program identifying multiple adjacent points as multiples of that BE, moving directly onto a search for the next BE.
3. *Flexibility* - that the length of time over which the next BE may occur can be defined within limits.

Thus, in the example above, recognition of BE_1 followed by BE_2 to give a LoCoD for one left stride followed by one right stride (Fig. 2) could be;

(BE_1) *Presence*	WHEN; $dAh_s/dt > 0.25\ g/s$ FOR $t > 0.2$ s IS TRUE
(BE_1) *Range*	SKIP DATA FOR 0.25 s
(BE_2) *Flexibility*	WHEN; $dAh_s/dt < -0.25\ g/s$ FOR $t > 0.2$ s WITHIN $t = 0.3$ s OF END OF BE_1

The value of the time-based definition is that it helps deal with variation in both amplitudes and periods of waveforms. Specifically, it allows the program to;

(a) be less susceptible to outliers (cf. *Presence*),

(b) detect the beginning of e.g. a waveform (cf. A1 in Fig. 2) and then allows flexibility in time to pass the peak of that waveform (cf. *Range*)

(c) constrain the length of time within which the next sub-element must occur for it to be considered part of the LoCoD (cf. *Flexibility*).

We present the computational process by which the data are treated using the LoCoD method in the supplementary material 1 but also note the following link (http://ggluck.swan.ac.uk/ftp/DDMT%20new%20version/) where the software can be downloaded.

Suggestions for defining Behavioural Elements

Although behavioural elements can be defined by simple inspection, the variability in the way they are manifest and the limits set to define them by the user are critical to the success of the overall algorithm for identifying behaviours. We suggest that the user first inspects the data in the form of line graphs over time to identify which data streams change predictably with the behaviour to be isolated. At this stage, the data can also be smoothed to reduce noise. In general, we note that running means are particularly valuable for smoothing out short-term outliers, diminishing noise and highlighting the major trends in waveforms; within the program above, the user can experiment with different smoothing windows to produce the clearest waveform in the data (cf. Fig. 1). Each data line to be used in the identification of a BE can then be cut from a number of examples of the BE in the data (ideally from a number of different animals) and these examples effectively superimposed on each other to show the variability in the data (Fig. 3). The same data can then be used to work out mean (and variance) numeric values for the parameters to be used in the (sensor value-based or time-based) rules to define a BE (Fig. 3). Consideration of the spread of the distribution of values of such parameters allows users to assess the extent to which the chosen thresholds will work within a population of the BEs.

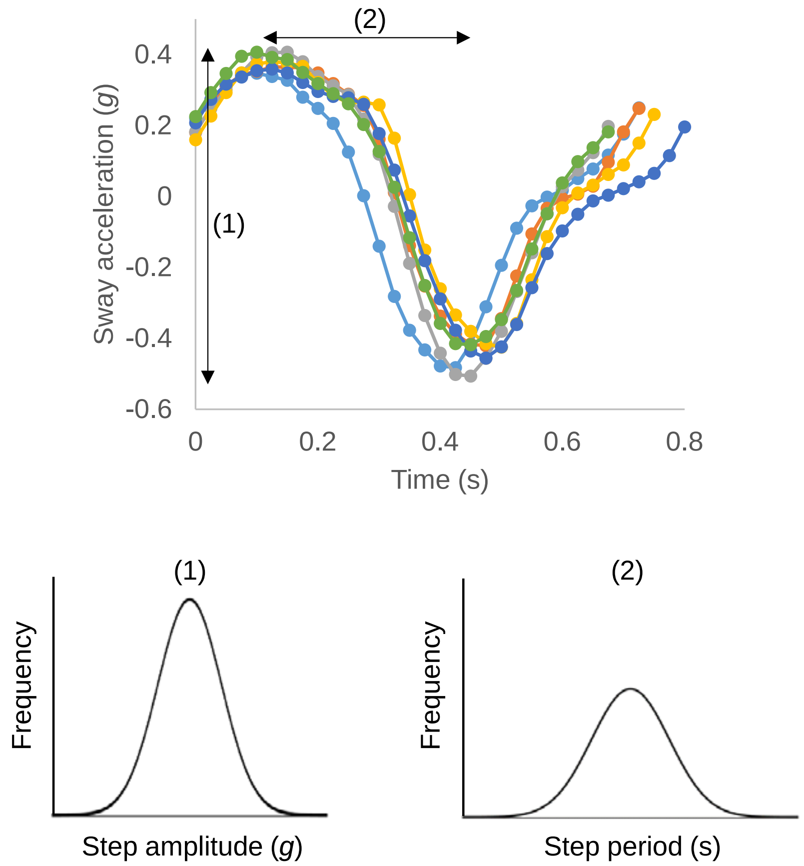

Fig. 3 – Example of the process of defining the value of parameters used to identify behavioural elements in the LoCoD method. The upper graph shows multiple examples of a given behaviour (penguin walking) in a recorded data stream that represents the behaviour well (in this case the smoothed sway acceleration). The superimposition of multiple examples of the behaviour highlights the variation in the behaviour. Construction of frequency distributions of particular elements that could be used to define a behavioural element (here step amplitude (1) and step period (2)) provide information on the probabilities of any given step falling outside user-defined limits to that distribution. This ultimately defines the extent to which the criteria will encompass the defined behavioural element.

In order to test the applicability of the LoCoD method over different behavioural periods, we used animal data corresponding to;

(a) ‘Short-period’ LoCoDs of behaviours, manifested by actions typically lasting less than 1 s: The examples used for this study were single bites of sheep and single steps by penguins walking.

(b) ‘Medium-period’ LoCoD of behaviours, typically lasting several seconds: The example used here was condors thermalling.

(c) ‘Long-period’ LoCoD of behaviours, typically lasting from between 30 s up to minutes: Here, we used cheetahs stalking prey.

A training dataset was created for behavioural identification for each of the above species, where all cases of the given behaviour was identified either according to known instances where the behaviour had been directly observed, or recorded, or by manual identification by an expert (see Supplementary material for behaviour descriptions and LoCoD definitions).

The LoCoD method was compared to other methods, see below, by considering the following metrics to assess classification performance: (1) Processing time (in seconds), which is the time spent by our single computer (to ensure that processing capacity was the same for all tasks) to identify and classify behaviours within defined data sets, and (2) Confusion Matrix-based scores: These metrics include Recall and Precision, which are routinely used in such comparisons (Resheff *et al.* 2014). Recall (also known as Sensitivity or True Positive rate) is estimated as: True Positives / (True Positives + False Negatives); and Precision is estimated as: True Positives / (True Positives + False Negatives). These two metrics are interesting because when Recall values increase, Precision values decrease, and we can assess the performance of a model by focusing the balance between both measures. By calculating both, we have a measure that expresses the ability of the model to find a particular behaviour in the dataset (i.e., Recall) while we have also a measure that expresses the proportion of the data points that our model classified as a particular behaviour that actually was that behaviour (i.e., Precision). We do not present Accuracy values for two reasons; i) since the LoCoD method does not consider each data point individually, quantification of the identification result of a given LoCoD case cannot give a true negative result and; ii) although true negative results can be established with the machine-learning methods, Accuracy can give biased results for unbalanced data sets (i.e., when the number of true positives in the confusion matrix is very different to the true negative (Sokolova & Lapalme 2009; Stąpor 2017).

2.2 │Comparator methods

We compared the outputs of the LoCoD method with nine different behavioural classifier models. These were; (1) K-Nearest Neighbours (K=3), (2) Linear Support Vector machines

(Linear SVM), (3) Radial Basis Function kernels for Support Vector Machines (RBF SVM), (4) Decision Trees, (5) Random Forest, (6) Naïve Bayes, (7) Linear Discriminant Analysis (LDA), (8) Quadratic Discriminant Analysis (QDA), (9) Artificial Neural Networks (ANN). These are all offered within a single piece of software as freeware (AccelerRater, http://accapp.move-ecol-minerva.huji.ac.il/) (Resheff *et al.* 2014) which facilitates protocols and testing (see a brief description of each model in Supplementary material). When using AccelerRater, we used the 'all features' option to construct the models (selecting the *"precomputed stats, Label"* option from the upload tab, to ensure that we could have available the same features employed with LoCoD) and a Train-Test split (50% for training and 50% for testing) for validation as for the LoCoD method. We note though, that machine learning methods have numerous options for fine tuning, which can have an appreciable impact on the overall accuracy (Ladds *et al.* 2017) so our comparison between machine learning options and the LoCoD method may have disadvantaged the machine learning process.

3 | RESULTS

The overall capacity of the LoCoD method to detect specified behaviours within varied datasets from free-living animals, was comparable, and sometimes higher, to some of the best methods otherwise tested (Tables 1 and 2). However, the speed with which the LoCoD method resolved behaviours was many times faster than the more conventional methods. For instance, the time required for the LoCoD method to process sheep biting and condor thermalling was less than 1% of the time required for the best machine-learning algorithm (representing 0.04% and 0.41%, respectively). In the case of the cheetah and the penguin data, the time required for the LoCoD method to classify the walking represented 6% and 20% of the total time required for the best machine-learning algorithm (Tables 1 & 2).

For sheep biting, although the best machine-learning algorithm (considering shortest processing time, together with highest recall and precision) was the QDA method, none of the used machine-learning algorithms had a good overall performance for classification (Table 1). The LoCoD method was the only approach that showed good performance in all the Confusion Matrix based scores (all above 85%).

For penguin walking, there were 4 machine-learning algorithms that showed similar performance for all metrics (Nearest Neighbour, RBF SVM, Decision Tree, and Random Forest).

The LoCoD method showed similar performance (Recall and Precision above 95%) but with processing times that were a fraction of the best machine-learning approaches (Table 1).

Although the best machine-learning method to classify condor thermalling was QDA, most of the methods resulted in poor performance, with most requiring excessive processing times and some even unable to provide a result (marked as NA in Table 2). The LoCoD method showed comparable performance to QDA, with lower Recall, higher Precision and notably lower processing times, equating to about 0.4% that of the QDA (Table 2). Although markedly slower than the LoCoD method (it took almost 250 times longer), the manual method outperformed all other options by an extended margin (Table 2).

In a manner similar to condor thermalling, most of the methods attempting to define cheetah stalking resulted in poor performance, many of them requiring excessive processing time, with the software from some systems unable to provide a result (marked as NA in Table 2). The best machine-learning method was Decision Tree. The LoCoD method showed comparable performance to this, with an approximately 10% lower Recall and Precision, but with significantly lower processing times, equating to about 6% that of the Decision Tree method (Table 2).

Overall, and of particular note, was that the LoCoD method dealt particularly well with behavioural identification where the temporal variability of the behaviour was high (defined by the range in duration of the different base elements of the behaviour). For example, in the case of the condor thermalling, manual labelling showed that each complete turn had a mean duration of 19.7 ± 4.9 s (SD), showing the variation in the presence, range and flexibility (cf. Fig. 3) of the two base elements used to define this behaviour (based on altitude gain and rates of change of magnetometry data - Supplementary Data, Table S3.3). Given that the sum of these three values limits the maximum duration of the LoCoD, all but one of the labelled LoCoD complete turns in thermal soaring were 15 seconds in duration. Similarly, where the machine-learning methods struggled with identification of the cheetah stalking, the LoCoD method performed well; the temporal range of this behaviour being 48.3 ± 16.2 s.

TABLE 1 Performance and time taken for the different identification methods to identify all cases of the ‘short-period’ behaviour of sheep biting and penguin walking in their respective data sets (see supplementary material for further detail). For each method, the time taken for the algorithm to run through the complete data set is given, along with the measures of recall and precision.

	Sheep biting			*Penguin walking*		
Method	Time (s)	Performance		Time (s)	Performance	
		Recall	Precision		Recall	Precision
Manual	2039	1.00	1.00	2040	1.00	1.00
LoCoD	**1.5**	**0.89**	**0.87**	**14**	**0.98**	**0.98**
Nearest Neighbour	243	0.00	0.00	77	0.97	0.96
Linear SVM	3189	0.00	0.00	359	1.00	0.75
RBF SVM	253	0.00	0.00	79	0.94	0.97
Decision Tree	242	0.00	0.00	80	0.97	0.96
Random Forest	281	0.00	0.00	82	0.98	0.96
Naïve Bayes	317	0.00	0.00	75	0.99	0.76
LDA	264	0.00	0.00	74	0.99	0.76
QDA	353	0.99	0.01	77	0.76	0.71
ANN	3451	0.00	0.00	405	0.92	0.97

TABLE 2 Performance and time taken for the different identification methods to identify all cases of 'medium-period' behaviour, consisting of condor thermalling and the 'long-period' behaviour of cheetah stalking in their respective data sets (see supplementary material for further detail). For these two behaviours, a number of machine-learning methods were not run to completion due to some system error, generally after more of 20 hours of processing time (marked with NA). For each method, the time taken for the algorithm to run through the complete data set is given, along with the measures of recall, and precision.

	Condor thermalling			***Cheetah stalking***		
Method	Time (s)	Performance		Time (s)	Performance	
		Recall	Precision		Recall	Precision
Manual	2220	1.00	1.00	180	1.00	1.00
LoCoD	**9**	**0.87**	**0.73**	**7.2**	**0.89**	**0.89**
Nearest Neighbour	2182	0.14	0.26	4045	0.99	0.98
Linear SVM	NA	NA	NA	NA	NA	NA
RBF SVM	NA	NA	NA	NA	NA	NA
Decision Tree	2358	0.01	0.35	3470	0.99	0.99
Random Forest	2998	0.00	0.00	4217	1.00	0.98
Naïve Bayes	NA	NA	NA	3179	0.19	0.03
LDA	2152	0.01	0.01	3016	0.06	0.26
QDA	2157	0.54	0.91	NA	NA	NA
ANN	NA	NA	NA	NA	NA	NA

4 │DISCUSSION

4.1 │ Speed *versus* accessibility considerations in identifying behaviours

In his seminal work on behaviour, Tinbergen (Tinbergen 1960) defined behaviours by noting prescribed changes in animal movement over time. This approach gets to the heart of behaviour description and is one that should be accessible by those using animal-attached sensors, e.g. accelerometers, magnetometers and gyroscopes (Johnson & Tyack 2003), that record body postures and movement in its various forms over time. Indeed, the precision with which movement descriptors such as angular velocity and acceleration can be measured has catalysed many studies of animal behaviour by workers using such smart tags (Yoda *et al.* 1999). More information about the movement from multiple sensors, many of which measure tri-axially to cover the 3 space dimensions anyway (Johnson & Tyack 2003; Wilson, Shepard & Liebsch 2008), can lead to very comprehensive descriptions of movement (Yoda, Kohno & Naito 2004), something that can be further enhanced by converting primary movement data (such as acceleration) to additional derivatives (such as VeDBA (Qasem *et al.* 2012)). Interpretation of such diverse and complex data is not intuitive, which makes a good case for machine-learning since no specialised knowledge is required by users. Coupled with this is the expectation that machine-learning systems produce best classifications if they are provided with most data, which makes a clear case for using all possible data (Resheff *et al.* 2014). However, this brings with it appreciable computational challenges because every new line of information has to be considered computationally with respect to all others. Processing time therefore increases disproportionately with the inclusion of every new data stream (Murphy 2012). Indeed, although computer processing speed continues to increase roughly according to Moore's Law, so too does our capacity to log data (Schaller 1997). Our ability to incorporate new sensors within our animal-attached tag systems (Ropert-Coudert & Wilson 2005), coupled with a proclivity to record at ever faster rates (Robert-Coudert & Wilson 2004) and derive new metrics from the base data (e.g. jerk, static- and dynamic acceleration as well as dynamic body acceleration from raw tri-axial acceleration data (Ydesen *et al.* 2014)) in tandem with tag deployments that may span months bringing in billions of prime data points, inevitably leads to more extended computing times.

Such a compromise might be more acceptable if the performance of machine-learning approaches was exceptional, but our results show that this is not the case (Tables 1 & 2). Our LoCoD approach requires good understanding and careful inspection of the sensor channels in order to make decisions about which data streams are most useful (and in which combination) to define the behaviour. This therefore requires some degree of specialist knowledge of the sensors used and an appreciable initial investment in time, although we would advocate that any use of sensor-acquired data 'blind' is not good practice anyway. Our suggestion is that the LoCoD approach specifically follows a 3 stage process; (1) where the primary data streams of interest are signal-processed to reduce noise and highlight patterns (e.g. via smoothing) over various scales, (2) where derived data streams, most notably differentials, are calculated for inclusion, if relevant (based on expectations and inspection of the behaviour in question) and (3) where conditions for sequential BEs are defined based on precise patterns in selected data streams with defined time-dependent flexibility for their execution. Such an approach is obviously more onerous for the worker than a machine-based learning technique and may be considered a disadvantage. However, this approach frees up appreciable amounts of computational time (Tables 1 & 2) by directing the machine to deal rapidly with a small fraction of the available data. This is critical for complex behaviours made up of many BEs. In the process, it allows identification of the minutia of behaviour if needed (e.g. left footsteps rather than 'walking') which may be important for rare, very short-lived behaviours. Indeed, the LoCoD method specifically identifies the smallest common denominator that defines a behaviour according to the sequence of BEs, for example single steps, or pairs of steps, within walking, rather than general walking *per se.* This leads to apparent overkill in that the approach will essentially identify every step during the tagged period, which may be more detail that many need, but steps within a defined time interval of each other can be merged without problem to produce larger bouts of walking if preferred and analysed according to behavioural type. Conversely, identification of slow, single steps, such as occur when herbivores graze, can lead to appreciable displacement over time, so their identification can be important in dead-reckoning approaches for resolving animal movement (Bidder *et al.* 2015). In addition, the ability to separate, for example, 'grazing and walking' from 'grazing without walking' should allow workers to recognise sub-behaviours within behaviours, something that is considered by people observing animals (Beker *et al.* 2010) but which are normally overlooked in tag data (Martiskainen *et al.* 2009). The LoCoD method performed slightly less well with our example of long-period behaviours than with short- or medium-period behaviours (Tables 1 & 2) making it apparently less useful (although the behaviour was identified in <0.5% of the time taken for the manual or machine-learning approach). Ultimately though, the absolute value of the approach depends on the extent to which the variability of the behaviour can be described by the flexibility of the algorithm used (see above). More work will be needed to determine the extent to which our results for cheetahs stalking are typical of 'long-lived' behaviours.

4.2 | Libraries of behaviours and inter-specific interpolation

An obvious advantage of explicitly defining an algorithm for a particular behaviour is that it can be stored and used for different individuals (cf. Fig. 3). However, a particular strength of the process of defining LoCoDs via BEs extends beyond this. This is because algorithms can be compared inter-specifically, and cognisance taken of changing values within the individual BEs to help predict what might be expected for new species. For example, the details of locomotion are known to be a broad function of mammal size and leg length (Christiansen 2002) so BEs coding for this should change in their specified conditions accordingly. Indeed, such specified conditions could be regressed against e.g. body mass to make predictions. As part of this general process, we anticipate that an online library could be created, which provides effective algorithms for determination of defined behaviours, which workers may readily consult for their own applications. Success in this venture may result in researchers using such algorithms without particular comprehension or time invested so that user expertise might eventually mirror those that employ machine-learning techniques. Against this, inter-specific variation beyond simple allometric expectations may serve to reduce the performance of this proposed cross-species approach (see Campbell *et al.* 2013). Either way though, having access to a defined method of determining the BEs within LoCoDs for behaviours for one species should certainly serve as a very useful starting point for users wishing to examine the same behaviour in another.

5 │Conclusions

Although the LoCoD method described here requires appreciable investment in time and understanding for workers to be able to develop appropriate algorithms for BEs, the approach clearly has value for those wishing to extract behaviours from multi-sensor data. The approach does not require a fixed sliding time window to operate, but has built-in flexibility in both time and amplitude to recognise patterns and, in addition, can be made to be 'blind' for a period within BEs so as not to be confused by the vagaries of variability at certain points within waveforms. This flexible template tactic, which uses a Boolean approach on only the bare minimum of data needed to recognise behaviours (ranging from those lasting less than 1 second to minutes or even hours, (cf. Horie *et al.* 2017), frees up processing time, making the whole process substantially more efficient. We would hope that algorithms for defined behaviours from particular species will be shared within the community to build up a potent library for the benefit of all wishing to try the approach.

ACKNOWLEDGEMENTS

This research contributes to the CAASE project funded by King Abdullah University of Science and Technology (KAUST) under the KAUST Sensor Initiative. We thank SANParks and the Department of Wildlife and National Parks, Botswana for allowing our research in the Kgalagadi Transfrontier Park (Permit Number 2006-05-01 MGLM) and Mr H. McKay and Mr W. Wright for their co-operation. We also thank the Nature Conservancy for allowing the sheep work in Patagonia. The study was made possible by various grants; a Leverhulme early career fellowship (ELCS), Fundación BBVa, PICT_BID-2014-0725 (SAL), the National Geographic Global Exploration Fund (#GEFNE89-13) the Royal Society (2009/R3 JP090604), NERC (NE/I002030/1) (SMS), and the Royal Society/Wolfson Lab refurbishment scheme (RPW). We also thank the Lewis Foundation, South Africa, The Howard G. Buffet Foundation, National Geographic, Kanabo Conservation Link, Comanis Foundation, Panthera and the Kruger Park Marathon Club for financial support. Hannah Williams thanks the Swansea University College of Science for their continued support and we thank Krüger Werke GmbH for technical help with waterproofing.

AUTHORS' CONTRIBUTIONS

RPW and MH conceived and developed the methodology, with input from all other authors. AdV, ELCS, FQ, JES, DMS and SL collected the data. AdV and HW tested various data sets with the algorithms and various machine-learning software and all authors contributed critically to the development of, and writing, the manuscript.

DATA ACCESSIBILITY

The data used in this work are deposited within Dryad.

ORCID

Rory Wilson (ID) http://orcid.org/0000-0003-3177-0177

REFERENCES

Altmann, J. (1974) Observational study of behavior: sampling methods. *Behaviour,* **49,** 227-266.
Anderberg, M.R. (2014) *Cluster analysis for applications: probability and mathematical statistics: a series of monographs and textbooks*. Academic press.
Beker, A., Gipson, T., Puchala, R., Askar, A., Tesfai, K., Detweiler, G., Asmare, A. & Goetsch, A. (2010) Energy expenditure and activity of different types of small ruminants grazing varying pastures in the summer. *Journal of Applied Animal Research,* **37,** 1-14.
Bidder, O., Walker, J., Jones, M., Holton, M., Urge, P., Scantlebury, D., Marks, N., Magowan, E., Maguire, I. & Wilson, R. (2015) Step by step: reconstruction of terrestrial animal movement paths by dead-reckoning. *Movement ecology,* **3,** 23.
Bidder, O.R., Campbell, H.A., Gómez-Laich, A., Urgé, P., Walker, J., Cai, Y., Gao, L., Quintana, F. & Wilson, R.P. (2014) Love thy neighbour: automatic animal behavioural classification of acceleration data using the k-nearest neighbour algorithm. *PloS one,* **9,** e88609.
Birkhead, T.R., Atkin, L. & Møller, A. (1987) Copulation behaviour of birds. *Behaviour,* **101,** 101-138.
Brown, D.D., Kays, R., Wikelski, M., Wilson, R. & Klimley, A.P. (2013) Observing the unwatchable through acceleration logging of animal behavior. *Animal Biotelemetry,* **1,** 1.
Campbell, H.A., Gao, L., Bidder, O.R., Hunter, J. & Franklin, C.E. (2013) Creating a behavioural classification module for acceleration data: using a captive surrogate for difficult to observe species. *Journal of Experimental Biology,* **216,** 4501-4506.
Chastin, S. & Granat, M. (2010) Methods for objective measure, quantification and analysis of sedentary behaviour and inactivity. *Gait & posture,* **31,** 82-86.
Christiansen, P. (2002) Locomotion in terrestrial mammals: the influence of body mass, limb length and bone proportions on speed. *Zoological Journal of the Linnean Society,* **136,** 685-714.
de Weerd, N., van Langevelde, F., van Oeveren, H., Nolet, B.A., Kölzsch, A., Prins, H.H. & de Boer, W.F. (2015) Deriving animal behaviour from high-frequency GPS: tracking cows in open and forested habitat. *PloS one,* **10,** e0129030.
Drews, C. (1993) The concept and definition of dominance in animal behaviour. *Behaviour,* **125,** 283-313.
Fourati, H., Manamanni, N., Afilal, L. & Handrich, Y. (2011) Posture and body acceleration tracking by inertial and magnetic sensing: Application in behavioral analysis of free-ranging animals. *Biomedical signal processing and control,* **6,** 94-104.
Hooker, S.K., Biuw, M., McConnell, B.J., Miller, P.J. & Sparling, C.E. (2007) Bio-logging science: logging and relaying physical and biological data using animal-attached tags. *Deep Sea Research Part II: Topical Studies in Oceanography,* **54,** 177-182.
Horie, J., Mitamura, H., Ina, Y., Mashino, Y., Noda, T., Moriya, K., Arai, N. & Sasakura, T. (2017) Development of a method for classifying and transmitting high-resolution feeding behavior of fish using an acceleration pinger. *Animal Biotelemetry,* **5,** 12.
Johnson, M.P. & Tyack, P.L. (2003) A digital acoustic recording tag for measuring the response of wild marine mammals to sound. *IEEE journal of oceanic engineering,* **28,** 3-12.
Krebs, J.R. & Davies, N.B. (2009) *Behavioural ecology: an evolutionary approach*. John Wiley & Sons.
Krebs, J.R., Davies, N.B. & Parr, J. (1993) *An introduction to behavioural ecology*. Blackwell Scientific Publications.
Ladds, M.A., Thompson, A.P., Kadar, J.-P., Slip, D., Hocking, D. & Harcourt, R. (2017) Super machine learning: improving accuracy and reducing variance of behaviour classification from accelerometry. *Animal Biotelemetry,* **5,** 8.
Laich, A.G., Wilson, R.P., Quintana, F. & Shepard, E.L. (2008) Identification of imperial cormorant Phalacrocorax atriceps behaviour using accelerometers. *Endangered Species Research,* **10,** 29-37.
Lucas, M.C. & Baras, E. (2000) Methods for studying spatial behaviour of freshwater fishes in the natural environment. *Fish and fisheries,* **1,** 283-316.
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.
Miller, N. & Gerlai, R. (2007) Quantification of shoaling behaviour in zebrafish (Danio rerio). *Behavioural brain research,* **184,** 157-166.
Moe-Nilssen, R. & Helbostad, J.L. (2004) Estimation of gait cycle characteristics by trunk accelerometry. *Journal of biomechanics,* **37,** 121-126.

Murphy, K.P. (2012) Machine Learning: A Probabilistic Perspective. Adaptive Computation and Machine Learning. MIT press.

Nathan, R., Spiegel, O., Fortmann-Roe, S., Harel, R., Wikelski, M. & Getz, W.M. (2012) Using tri-axial acceleration data to identify behavioral modes of free-ranging animals: general concepts and tools illustrated for griffon vultures. *Journal of Experimental Biology,* **215,** 986-996.

Noda, T., Kawabata, Y., Arai, N., Mitamura, H. & Watanabe, S. (2014) Animal-mounted gyroscope/accelerometer/magnetometer: In situ measurement of the movement performance of fast-start behaviour in fish. *Journal of Experimental Marine Biology and Ecology,* **451,** 55-68.

Qasem, L., Cardew, A., Wilson, A., Griffiths, I., Halsey, L.G., Shepard, E.L., Gleiss, A.C. & Wilson, R. (2012) Tri-axial dynamic acceleration as a proxy for animal energy expenditure; should we be summing values or calculating the vector? *PLoS One,* **7,** e31187.

Resheff, Y.S., Rotics, S., Harel, R., Spiegel, O. & Nathan, R. (2014) AcceleRater: a web application for supervised learning of behavioral modes from acceleration measurements. *Movement ecology,* **2,** 27.

Robert-Coudert, Y. & Wilson, R.P. (2004) Subjectivity in bio-logging: do logged data mislead? *Memoirs of the National Institute of Polar Research: Special Issue,* **58,** 23-33.

Rong, L., Zhiguo, D., Jianzhong, Z. & Ming, L. (2007) Identification of individual walking patterns using gait acceleration. *Bioinformatics and Biomedical Engineering, 2007. ICBBE 2007. The 1st International Conference on*, pp. 543-546. IEEE.

Ropert-Coudert, Y. & Wilson, R.P. (2005) Trends and perspectives in animal ‐ attached remote sensing. *Frontiers in Ecology and the Environment,* **3,** 437-444.

Sabatini, A.M., Martelloni, C., Scapellato, S. & Cavallo, F. (2005) Assessment of walking features from foot inertial sensing. *IEEE Transactions on biomedical engineering,* **52,** 486-494.

Sakamoto, K.Q., Sato, K., Ishizuka, M., Watanuki, Y., Takahashi, A., Daunt, F. & Wanless, S. (2009) Can ethograms be automatically generated using body acceleration data from free-ranging birds? *PloS one,* **4,** e5379.

Samarasinghe, S. (2016) *Neural networks for applied sciences and engineering: from fundamentals to complex pattern recognition*. CRC Press.

Schaller, R.R. (1997) Moore's law: past, present and future. *IEEE spectrum,* **34,** 52-59.

Shepard, E.L., Wilson, R.P., Quintana, F., Laich, A.G., Liebsch, N., Albareda, D.A., Halsey, L.G., Gleiss, A., Morgan, D.T. & Myers, A.E. (2008) Identification of animal movement patterns using tri-axial accelerometry. *Endangered Species Research,* **10,** 47-60.

Sokolova, M. & Lapalme, G. (2009) A systematic analysis of performance measures for classification tasks. *Information Processing & Management,* **45,** 427-437.

Stąpor, K. (2017) Evaluating and Comparing Classifiers: Review, Some Recommendations and Limitations. *International Conference on Computer Recognition Systems*, pp. 12-21. Springer.

Tachibana, R.O., Oosugi, N. & Okanoya, K. (2014) Semi-automatic classification of birdsong elements using a linear support vector machine. *PloS one,* **9,** e92584.

Tinbergen, N. (1960) Comparative Studies of the Behaviour of Gulls (Laridae): a Progress Report1. *Behaviour,* **15,** 1-69.

Walker, J.S., Jones, M.W., Laramee, R.S., Bidder, O.R., Williams, H.J., Scott, R., Shepard, E.L. & Wilson, R.P. (2015a) TimeClassifier: a visual analytic system for the classification of multi-dimensional time series data. *The Visual Computer,* **31,** 1067-1078.

Walker, J.S., Jones, M.W., Laramee, R.S., Holton, M.D., Shepard, E.L., Williams, H.J., Scantlebury, D.M., Nikki, J.M., Magowan, E.A. & Maguire, I.E. (2015b) Prying into the intimate secrets of animal lives; software beyond hardware for comprehensive annotation in ‘Daily Diary’tags. *Movement ecology,* **3,** 1.

Williams, H.J., Holton, M.D., Shepard, E.L., Largey, N., Norman, B., Ryan, P.G., Duriez, O., Scantlebury, M., Quintana, F. & Magowan, E.A. (2017) Identification of animal movement patterns using tri-axial magnetometry. *Movement ecology,* **5,** 6.

Wilson, R., Culik, B., Bannasch, R. & Lage, J. (1994) Monitoring Antarctic environmental variables using penguins. *Marine Ecology Progress Series,* **106,** 199-202.

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.

Ydesen, K.S., Wisniewska, D.M., Hansen, J.D., Beedholm, K., Johnson, M. & Madsen, P.T. (2014) What a jerk: prey engulfment revealed by high-rate, super-cranial accelerometry on a harbour seal (Phoca vitulina). *Journal of Experimental Biology,* **217,** 2239-2243.

Yoda, K., Kohno, H. & Naito, Y. (2004) Development of flight performance in the brown booby. *Proceedings of the Royal Society of London B: Biological Sciences,* **271,** S240-S242.

Yoda, K., Sato, K., Niizuma, Y., Kurita, M., Bost, C., Le Maho, Y. & Naito, Y. (1999) Precise monitoring of porpoising behaviour of Adélie penguins determined using acceleration data loggers. *Journal of Experimental Biology,* **202,** 3121-3126.