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

33 become more likely to become processor limited and thereby take appreciable amounts

34 of time to resolve behaviours.

35 3. We suggest a Boolean approach whereby critical changes in recorded parameters are

36 used as sequential templates with defined flexibility (in both time and degree) to

37 determine individual behavioural elements within a behavioural sequence that, together,

38 makes up a single, defined behaviour.

39 4. We tested this approach, and compared it to a suite of other behavioural identification

40 methods, on a number of behaviours from tag-equipped animals; sheep grazing, penguins

41 walking, cheetah stalking prey and condors thermalling.

42 5. Overall behaviour recognition using our new approach was better than most other

43 methods due to; (i) its ability to deal with behavioural variation and (ii) the speed with

44 which the task was completed because extraneous data are avoided in the process.

45 6. We suggest that this approach is a promising way forward in an increasingly data-rich

46 environment and that workers sharing algorithms can provide a powerful library for the

47 benefit of all involved in such work.

48

49 **1 │ INTRODUCTION**

50

51 Animal behaviour has been variously defined, but generally can be defined as 'the way in which

52 an animal works, functions or responds to a particular situation' (Tinbergen 1960) with

53 consequences for lifetime reproductive success (Birkhead, Atkin & Møller 1987; Drews 1993;

54 Krebs, Davies & Parr 1993; Krebs & Davies 2009). As such, our ability to determine animal

55 behaviours precisely is critically important for proper understanding of animal ecology and

56 ecosystem functioning (Krebs, Davies & Parr 1993). Indeed, it is this that explains the large

57 variety of methodologies developed to quantify behaviour (e.g. Tinbergen 1960; Altmann 1974;

58 Lucas & Baras 2000; Miller & Gerlai 2007; Chastin & Granat 2010). A particularly rapidly

59 developing field in this regard is 'biologging' – the deployment of autonomous tags on animals

60 to record data (Hooker *et al.* 2007). Specifically, the extraordinary development of electronic

61 technology over the last 3 decades has led the progression of sophisticated miniature sensors

62 coupled with low power consumption and rapidly expanding memory capacity (Ropert-Coudert

63 & Wilson 2005) so that studies using multi-sensor technology in tags on animals are now

64 common (Brown *et al.* 2013). This has led from the simple animal-attached tags of the 1990s

65 recording data once every few seconds (Wilson *et al.* 1994), to systems today that may record

66   multiple channels at thousands of Hertz (Johnson & Tyack 2003). Of particular note for defining

67   behaviours is the role played by accelerometers, gyroscopes and magnetometers, which can

68   resolve both animal attitude in the 3 spatial axes (Yoda *et al.* 1999; Williams *et al.* 2017) and

69   movement (Fourati *et al.* 2011; Noda *et al.* 2014). These are primary elements used in classifying

70   behaviours (Tinbergen 1960), and so have great potential in studies of wild animals.

71   However, the ease with which we can now record the physical manifestation of

72   behaviour, *via* metrics such as pitch, roll and 'dynamism' in the acceleration signature (Laich *et*

73   *al.* 2008), is tempered by the difficulties of dealing with the complexity and volume of such data.

74   Thus, computational solutions for processing the signals are inevitable and, accordingly, there is

75   a rich and varied literature dealing with this (e.g. Sakamoto *et al.* 2009; Nathan *et al.* 2012;

76   Resheff *et al.* 2014). This includes support vector machines (Tachibana, Oosugi & Okanoya

77   2014), regression trees (de Weerd *et al.* 2015), random forests (Bidder *et al.* 2014), neural

78   networks (Samarasinghe 2016), linear discriminant analysis (Anderberg 2014) and template-

79   matching (Walker *et al.* 2015b). Each method has advantages and disadvantages (Resheff *et al.*

80   2014) but prime negative issues revolve around subjectivity, whether the data are parametric, the

81   extent of over-fitting, and the computational time involved in the process (Nathan *et al.* 2012).

82   In addition, a particular weakness of many systems is that they fail to recognise the temporal

83   sequencing of the movements that define the fundamental unit of that behaviour and the

84   variability within them, and thereby preclude an important discriminator. For example, walking

85   may be defined by a cluster of acceleration metrics (Bidder *et al.* 2014) but the fundamental unit

86   of walking is the single step (Moe-Nilssen & Helbostad 2004) and this has well-defined

87   properties over time (Sabatini *et al.* 2005) that could, for example, be used in any decision tree-

88   based approach.

89   In this paper, we present an approach for identifying behaviours from data derived from

90   animal-attached tags that recognises (i) the lowest common denominator (LoCoD) defining any

91   particular behaviour (i.e. a single step is the lowest common denominator within walking) and

92   (ii) that this lowest common denominator can be usefully broken down into base elements (BEs)

93   (such as an increase, followed by a drop, in dorso-ventral acceleration for walking (Rong *et al.*

94   2007)), all of which have to follow each other in a defined sequence for the LoCoD to be

95   apparent. Finally, (iii), the timing of BEs within a sequence is often constrained. Thus, this

96   process provides a recognizable key for LoCoDs of behaviours based on measurements,

97   sequences and timings of BEs. We appreciate that much of the essence of this is inherent in some

98   template-matching approaches (Walker *et al.* 2015a) but combine this with both temporal

99 flexibility across all BEs, together with an ability to switch between and incorporate defined,
100 often derived, metrics that provide critical information for a powerful match. We demonstrate
101 the utility of this approach by using it to search for behaviours that have LoCoD periods ranging
102 between fractions of a second and several minutes using data derived from animal-attached tags
103 and compare it briefly to other computational methods.

104

105 **2 │MATERIALS AND METHODS**

106

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

111

112 **2.1 │ The LoCoD Method**

113

114 The LoCoD method involves initial consideration of the data visually by the user, who should
115 examine the details of the movement that makes up the behaviour and reflect how this movement
116 is expected to affect the sensors. In this, the user should identify the patterns that make up the
117 BEs of the LoCoD and whether they can be made more distinctive by selective smoothing, as is
118 done in many behaviour-identifying protocols anyway (Nathan *et al.* 2012). In addition, it is
119 recommended that differentials be derived for any signals of interest, since these often act as
120 excellent thresholds in derivation of the BEs (Fig. 1). Differentials are particularly important
121 since postural data derived from acceleration (Shepard *et al.* 2008) are dependent, in part, on the
122 angle of the terrain beneath the study animal (cf. the difference in sway axis during the stationary
123 periods at the beginning and end of the walking period in Fig. 1), as well as the tag placement.
124 Thus, working with differentials essentially standardises the signal output.
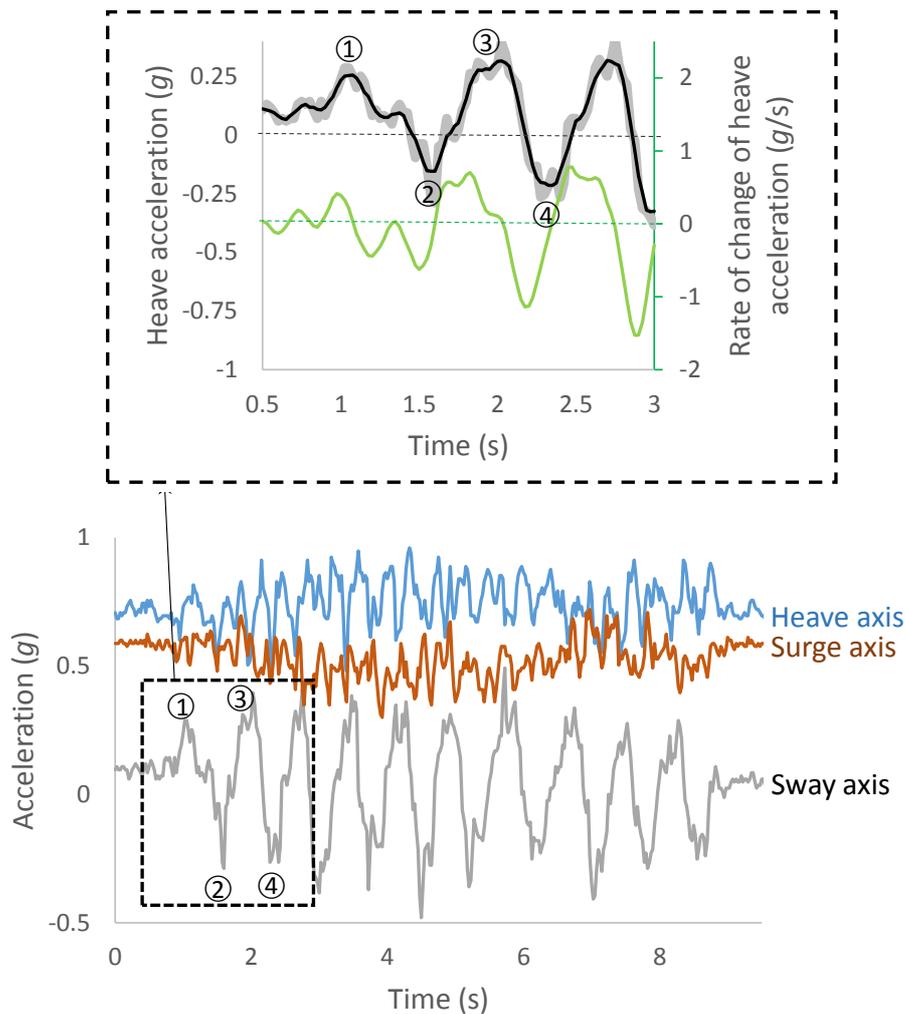
125

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 - RECOGNISE WHEN; $dAh_s/dt > 0.25$ g/s   (1)

where Ah_s is the smoothed heave acceleration following;

136

137 $$Ahs = \frac{1}{n} \sum_{i=0}^{n-1} Ah - i \qquad (2)$$

138
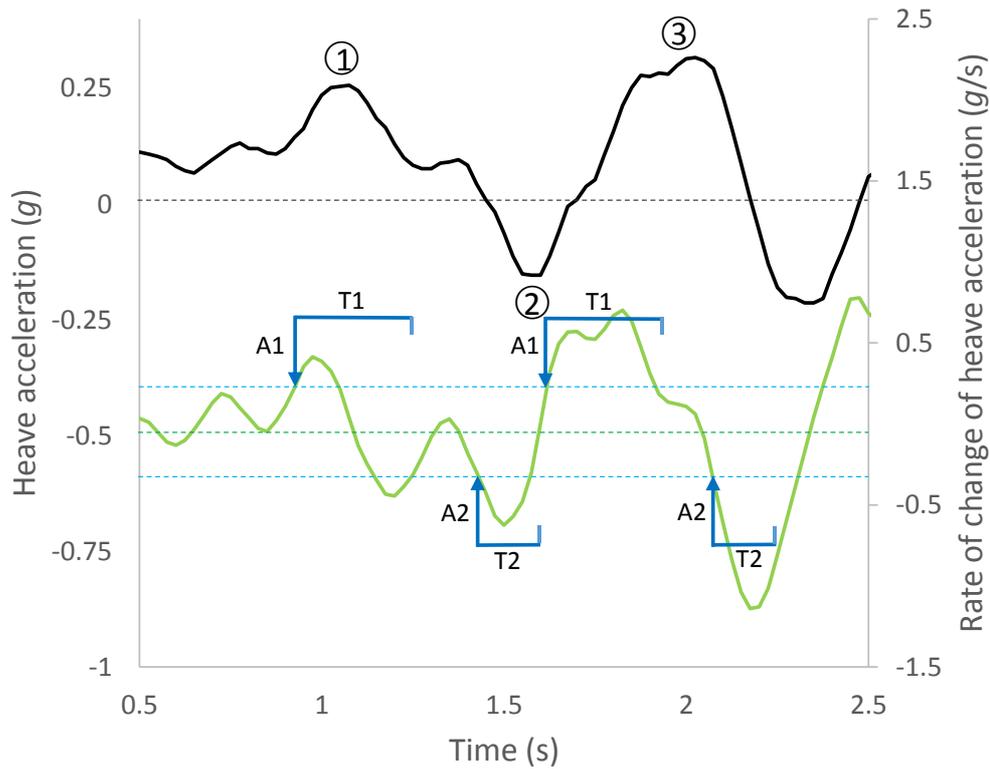
Fig. 2 – The first 3 steps (numbered) of the walking period shown in Fig. 1 for the smoothed heave axis (black line) and the rate of change of the heave axis (green line). The LoCoD method first identifies a feature, or combination of features, that signify the initiation of the first BE of the behaviour (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.

139

140

141 In addition, the process should recognise multiple, cross-channel sub-conditions (for positives

142 and negatives). Thus, equation (1) might be made of 3 sub-conditions;

143

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

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

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

147

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

149

150 Importantly, each sub-condition or condition can employ a time base with three elements within

151 it that can be specified. These are;

152

153   1. *Presence* - that the sub-condition or condition is maintained over a specified time for the

154   statement to be TRUE

155   2. *Range* - that, following identification of a true sub-condition or condition, the program

156   can skip a defined number of data points before looking for the next BE. This is important

157   because it can stop the program identifying multiple adjacent points as multiples of that

158   BE, moving directly onto a search for the next BE.

159   3. *Flexibility* - that the length of time over which the next BE may occur can be defined

160   within limits.

161

162 Thus, in the example above, recognition of BE_1 followed by BE_2 to give a LoCoD for one left

163 stride followed by one right stride (Fig. 2) could be;

164

165 (BE_1) *Presence*   WHEN; $dAh_s/dt > 0.25$ g/s FOR $t > 0.2$ s IS TRUE

166 (BE_1) *Range*   SKIP DATA FOR 0.25 s

167 (BE_2) *Flexibility*   WHEN; $dAh_s/dt < -0.25$ g/s FOR $t > 0.2$ s WITHIN $t = 0.3$ s OF END OF

168   BE_1

169

170 The value of the time-based definition is that it helps deal with variation in both amplitudes and

171 periods of waveforms. Specifically, it allows the program to;

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

173   (b) detect the beginning of e.g. a waveform (cf. A1 in Fig. 2) and then allows flexibility in

174   time to pass the peak of that waveform (cf. *Range*)

175   (c) constrain the length of time within which the next sub-element must occur for it to be

176   considered part of the LoCoD (cf. *Flexibility*).

177 We present the computational process by which the data are treated using the LoCoD method

178 in the supplementary material 1 but also note the following link

179 (http://ggluck.swan.ac.uk/ftp/DDMT%20new%20version/) where the software can be downloaded.

180

181 *Suggestions for defining Behavioural Elements*

182 Although behavioural elements can be defined by simple inspection, the variability in the way

183 they are manifest and the limits set to define them by the user are critical to the success of the

184 overall algorithm for identifying behaviours. We suggest that the user first inspects the data in

185 the form of line graphs over time to identify which data streams change predictably with the

186 behaviour to be isolated. At this stage, the data can also be smoothed to reduce noise. In general,

187 we note that running means are particularly valuable for smoothing out short-term outliers,

188 diminishing noise and highlighting the major trends in waveforms; within the program above,

189 the user can experiment with different smoothing windows to produce the clearest waveform in

190 the data (cf. Fig. 1). Each data line to be used in the identification of a BE can then be cut from

191 a number of examples of the BE in the data (ideally from a number of different animals) and

192 these examples effectively superimposed on each other to show the variability in the data (Fig.

193 3). The same data can then be used to work out mean (and variance) numeric values for the

194 parameters to be used in the (sensor value-based or time-based) rules to define a BE (Fig. 3).

195 Consideration of the spread of the distribution of values of such parameters allows users to assess

196 the extent to which the chosen thresholds will work within a population of the BEs.
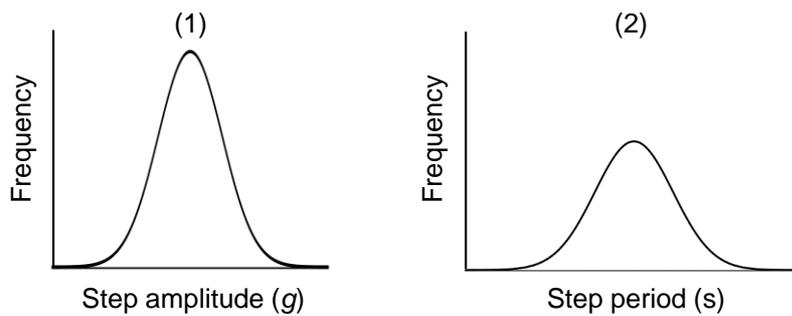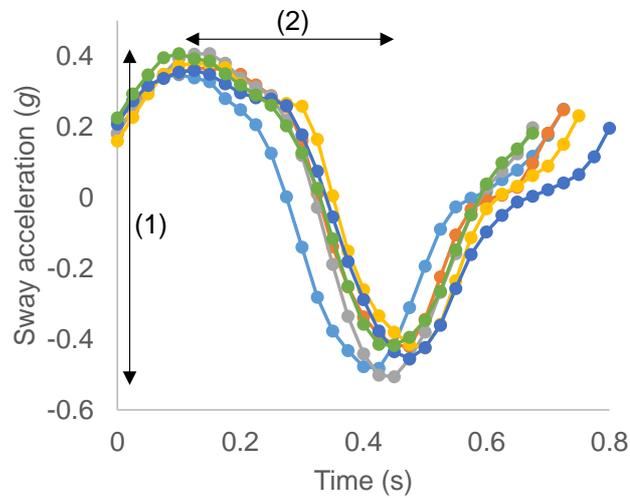
Fig. 3 – Example of the process of defining the value of parameters used to identify behavioural elements in the LoCoD method. The upper graph shows multiple examples of a given behaviour (penguin walking) in a recorded data stream that represents the behaviour well (in this case the smoothed sway acceleration). 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.

197

198 In order to test the applicability of the LoCoD method over different behavioural periods, we

199 used animal data corresponding to;

200

201 (a) 'Short-period' LoCoDs of behaviours, manifested by actions typically lasting less than 1

202   s: The examples used for this study were single bites of sheep and single steps by

203   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

237 (Linear SVM), (3) Radial Basis Function kernels for Support Vector Machines (RBF SVM), (4)

238 Decision Trees, (5) Random Forest, (6) Naïve Bayes, (7) Linear Discriminant Analysis (LDA),

239 (8) Quadratic Discriminant Analysis (QDA), (9) Artificial Neural Networks (ANN). These are

240 all offered within a single piece of software as freeware (AccelerRater, http://accapp.move-ecol-

241 minerva.huji.ac.il/) (Resheff *et al.* 2014) which facilitates protocols and testing (see a brief

242 description of each model in Supplementary material). When using AccelerRater, we used the

243 'all features' option to construct the models (selecting the *"precomputed stats, Label"* option from

244 the upload tab, to ensure that we could have available the same features employed with LoCoD)

245 and a Train-Test split (50% for training and 50% for testing) for validation as for the LoCoD

246 method. We note though, that machine learning methods have numerous options for fine tuning,

247 which can have an appreciable impact on the overall accuracy (Ladds *et al.* 2017) so our

248 comparison between machine learning options and the LoCoD method may have disadvantaged

249 the machine learning process.

250

251 **3 │ RESULTS**

252

253 The overall capacity of the LoCoD method to detect specified behaviours within varied datasets

254 from free-living animals, was comparable, and sometimes higher, to some of the best methods

255 otherwise tested (Tables 1 and 2). However, the speed with which the LoCoD method resolved

256 behaviours was many times faster than the more conventional methods. For instance, the time

257 required for the LoCoD method to process sheep biting and condor thermalling was less than

258 1% of the time required for the best machine-learning algorithm (representing 0.04% and 0.41%,

259 respectively). In the case of the cheetah and the penguin data, the time required for the LoCoD

260 method to classify the walking represented 6% and 20% of the total time required for the best

261 machine-learning algorithm (Tables 1 & 2).

262 For sheep biting, although the best machine-learning algorithm (considering shortest

263 processing time, together with highest recall and precision) was the QDA method, none of the

264 used machine-learning algorithms had a good overall performance for classification (Table 1).

265 The LoCoD method was the only approach that showed good performance in all the Confusion

266 Matrix based scores (all above 85%).

267 For penguin walking, there were 4 machine-learning algorithms that showed similar

268 performance for all metrics (Nearest Neighbour, RBF SVM, Decision Tree, and Random Forest).

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

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

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

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

296

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

Method	*Sheep biting*			*Penguin walking*		
	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

303

304

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	Performance		Time	Performance	
	(s)	Recall	Precision	(s)	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

321 description and is one that should be accessible by those using animal-attached sensors, e.g.
322 accelerometers, magnetometers and gyroscopes (Johnson & Tyack 2003), that record body
323 postures and movement in its various forms over time. Indeed, the precision with which
324 movement descriptors such as angular velocity and acceleration can be measured has catalysed
325 many studies of animal behaviour by workers using such smart tags (Yoda *et al.* 1999). More
326 information about the movement from multiple sensors, many of which measure tri-axially to
327 cover the 3 space dimensions anyway (Johnson & Tyack 2003; Wilson, Shepard & Liebsch
328 2008), can lead to very comprehensive descriptions of movement (Yoda, Kohno & Naito 2004),
329 something that can be further enhanced by converting primary movement data (such as
330 acceleration) to additional derivatives (such as VeDBA (Qasem *et al.* 2012)). Interpretation of
331 such diverse and complex data is not intuitive, which makes a good case for machine-learning
332 since no specialised knowledge is required by users. Coupled with this is the expectation that
333 machine-learning systems produce best classifications if they are provided with most data, which
334 makes a clear case for using all possible data (Resheff *et al.* 2014). However, this brings with it
335 appreciable computational challenges because every new line of information has to be
336 considered computationally with respect to all others. Processing time therefore increases
337 disproportionately with the inclusion of every new data stream (Murphy 2012). Indeed, although
338 computer processing speed continues to increase roughly according to Moore's Law, so too does
339 our capacity to log data (Schaller 1997). Our ability to incorporate new sensors within our
340 animal-attached tag systems (Ropert-Coudert & Wilson 2005), coupled with a proclivity to
341 record at ever faster rates (Robert-Coudert & Wilson 2004) and derive new metrics from the
342 base data (e.g. jerk, static- and dynamic acceleration as well as dynamic body acceleration from
343 raw tri-axial acceleration data (Ydesen *et al.* 2014)) in tandem with tag deployments that may
344 span months bringing in billions of prime data points, inevitably leads to more extended
345 computing times.

346 Such a compromise might be more acceptable if the performance of machine-learning
347 approaches was exceptional, but our results show that this is not the case (Tables 1 & 2). Our
348 LoCoD approach requires good understanding and careful inspection of the sensor channels in
349 order to make decisions about which data streams are most useful (and in which combination)
350 to define the behaviour. This therefore requires some degree of specialist knowledge of the
351 sensors used and an appreciable initial investment in time, although we would advocate that any
352 use of sensor-acquired data 'blind' is not good practice anyway. Our suggestion is that the
353 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

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

453

454   Altmann, J. (1974) Observational study of behavior: sampling methods. *Behaviour,* **49,** 227-266.
455   Anderberg, M.R. (2014) *Cluster analysis for applications: probability and mathematical statistics: a*
456   *series of monographs and textbooks*. Academic press.
457   Beker, A., Gipson, T., Puchala, R., Askar, A., Tesfai, K., Detweiler, G., Asmare, A. & Goetsch, A. (2010)
458   Energy expenditure and activity of different types of small ruminants grazing varying pastures
459   in the summer. *Journal of Applied Animal Research,* **37,** 1-14.
460   Bidder, O., Walker, J., Jones, M., Holton, M., Urge, P., Scantlebury, D., Marks, N., Magowan, E.,
461   Maguire, I. & Wilson, R. (2015) Step by step: reconstruction of terrestrial animal movement
462   paths by dead-reckoning. *Movement ecology,* **3,** 23.
463   Bidder, O.R., Campbell, H.A., Gómez-Laich, A., Urgé, P., Walker, J., Cai, Y., Gao, L., Quintana, F. &
464   Wilson, R.P. (2014) Love thy neighbour: automatic animal behavioural classification of
465   acceleration data using the k-nearest neighbour algorithm. *PloS one,* **9,** e88609.
466   Birkhead, T.R., Atkin, L. & Møller, A. (1987) Copulation behaviour of birds. *Behaviour,* **101,** 101-138.
467   Brown, D.D., Kays, R., Wikelski, M., Wilson, R. & Klimley, A.P. (2013) Observing the unwatchable
468   through acceleration logging of animal behavior. *Animal Biotelemetry,* **1,** 1.
469   Campbell, H.A., Gao, L., Bidder, O.R., Hunter, J. & Franklin, C.E. (2013) Creating a behavioural
470   classification module for acceleration data: using a captive surrogate for difficult to observe
471   species. *Journal of Experimental Biology,* **216,** 4501-4506.
472   Chastin, S. & Granat, M. (2010) Methods for objective measure, quantification and analysis of
473   sedentary behaviour and inactivity. *Gait & posture,* **31,** 82-86.
474   Christiansen, P. (2002) Locomotion in terrestrial mammals: the influence of body mass, limb length
475   and bone proportions on speed. *Zoological Journal of the Linnean Society,* **136,** 685-714.
476   de Weerd, N., van Langevelde, F., van Oeveren, H., Nolet, B.A., Kölzsch, A., Prins, H.H. & de Boer, W.F.
477   (2015) Deriving animal behaviour from high-frequency GPS: tracking cows in open and
478   forested habitat. *PloS one,* **10,** e0129030.
479   Drews, C. (1993) The concept and definition of dominance in animal behaviour. *Behaviour,* **125,** 283-
480   313.
481   Fourati, H., Manamanni, N., Afilal, L. & Handrich, Y. (2011) Posture and body acceleration tracking by
482   inertial and magnetic sensing: Application in behavioral analysis of free-ranging animals.
483   *Biomedical signal processing and control,* **6,** 94-104.
484   Hooker, S.K., Biuw, M., McConnell, B.J., Miller, P.J. & Sparling, C.E. (2007) Bio-logging science: logging
485   and relaying physical and biological data using animal-attached tags. *Deep Sea Research Part*
486   *II: Topical Studies in Oceanography,* **54,** 177-182.
487   Horie, J., Mitamura, H., Ina, Y., Mashino, Y., Noda, T., Moriya, K., Arai, N. & Sasakura, T. (2017)
488   Development of a method for classifying and transmitting high-resolution feeding behavior of
489   fish using an acceleration pinger. *Animal Biotelemetry,* **5,** 12.
490   Johnson, M.P. & Tyack, P.L. (2003) A digital acoustic recording tag for measuring the response of wild
491   marine mammals to sound. *IEEE journal of oceanic engineering,* **28,** 3-12.
492   Krebs, J.R. & Davies, N.B. (2009) *Behavioural ecology: an evolutionary approach*. John Wiley & Sons.
493   Krebs, J.R., Davies, N.B. & Parr, J. (1993) *An introduction to behavioural ecology*. Blackwell Scientific
494   Publications.
495   Ladds, M.A., Thompson, A.P., Kadar, J.-P., Slip, D., Hocking, D. & Harcourt, R. (2017) Super machine
496   learning: improving accuracy and reducing variance of behaviour classification from
497   accelerometry. *Animal Biotelemetry,* **5,** 8.
498   Laich, A.G., Wilson, R.P., Quintana, F. & Shepard, E.L. (2008) Identification of imperial cormorant
499   Phalacrocorax atriceps behaviour using accelerometers. *Endangered Species Research,* **10,** 29-
500   37.
501   Lucas, M.C. & Baras, E. (2000) Methods for studying spatial behaviour of freshwater fishes in the
502   natural environment. *Fish and fisheries,* **1,** 283-316.
503   Martiskainen, P., Järvinen, M., Skön, J.-P., Tiirikainen, J., Kolehmainen, M. & Mononen, J. (2009) Cow
504   behaviour pattern recognition using a three-dimensional accelerometer and support vector
505   machines. *Applied Animal Behaviour Science,* **119,** 32-38.
506   Miller, N. & Gerlai, R. (2007) Quantification of shoaling behaviour in zebrafish (Danio rerio).
507   *Behavioural brain research,* **184,** 157-166.
508   Moe-Nilssen, R. & Helbostad, J.L. (2004) Estimation of gait cycle characteristics by trunk
509   accelerometry. *Journal of biomechanics,* **37,** 121-126.

510 Murphy, K.P. (2012) Machine Learning: A Probabilistic Perspective. Adaptive Computation and
511   Machine Learning. MIT press.
512 Nathan, R., Spiegel, O., Fortmann-Roe, S., Harel, R., Wikelski, M. & Getz, W.M. (2012) Using tri-axial
513   acceleration data to identify behavioral modes of free-ranging animals: general concepts and
514   tools illustrated for griffon vultures. *Journal of Experimental Biology,* **215,** 986-996.
515 Noda, T., Kawabata, Y., Arai, N., Mitamura, H. & Watanabe, S. (2014) Animal-mounted
516   gyroscope/accelerometer/magnetometer: In situ measurement of the movement
517   performance of fast-start behaviour in fish. *Journal of Experimental Marine Biology and*
518   *Ecology,* **451,** 55-68.
519 Qasem, L., Cardew, A., Wilson, A., Griffiths, I., Halsey, L.G., Shepard, E.L., Gleiss, A.C. & Wilson, R.
520   (2012) Tri-axial dynamic acceleration as a proxy for animal energy expenditure; should we be
521   summing values or calculating the vector? *PLoS One,* **7,** e31187.
522 Resheff, Y.S., Rotics, S., Harel, R., Spiegel, O. & Nathan, R. (2014) AcceleRater: a web application for
523   supervised learning of behavioral modes from acceleration measurements. *Movement*
524   *ecology,* **2,** 27.
525 Robert-Coudert, Y. & Wilson, R.P. (2004) Subjectivity in bio-logging: do logged data mislead? *Memoirs*
526   *of the National Institute of Polar Research: Special Issue,* **58,** 23-33.
527 Rong, L., Zhiguo, D., Jianzhong, Z. & Ming, L. (2007) Identification of individual walking patterns using
528   gait acceleration. *Bioinformatics and Biomedical Engineering, 2007. ICBBE 2007. The 1st*
529   *International Conference on*, pp. 543-546. IEEE.
530 Ropert-Coudert, Y. & Wilson, R.P. (2005) Trends and perspectives in animal‐attached remote
531   sensing. *Frontiers in Ecology and the Environment,* **3,** 437-444.
532 Sabatini, A.M., Martelloni, C., Scapellato, S. & Cavallo, F. (2005) Assessment of walking features from
533   foot inertial sensing. *IEEE Transactions on biomedical engineering,* **52,** 486-494.
534 Sakamoto, K.Q., Sato, K., Ishizuka, M., Watanuki, Y., Takahashi, A., Daunt, F. & Wanless, S. (2009) Can
535   ethograms be automatically generated using body acceleration data from free-ranging birds?
536   *PloS one,* **4,** e5379.
537 Samarasinghe, S. (2016) *Neural networks for applied sciences and engineering: from fundamentals to*
538   *complex pattern recognition*. CRC Press.
539 Schaller, R.R. (1997) Moore's law: past, present and future. *IEEE spectrum,* **34,** 52-59.
540 Shepard, E.L., Wilson, R.P., Quintana, F., Laich, A.G., Liebsch, N., Albareda, D.A., Halsey, L.G., Gleiss, A.,
541   Morgan, D.T. & Myers, A.E. (2008) Identification of animal movement patterns using tri-axial
542   accelerometry. *Endangered Species Research,* **10,** 47-60.
543 Sokolova, M. & Lapalme, G. (2009) A systematic analysis of performance measures for classification
544   tasks. *Information Processing & Management,* **45,** 427-437.
545 Stąpor, K. (2017) Evaluating and Comparing Classifiers: Review, Some Recommendations and
546   Limitations. *International Conference on Computer Recognition Systems*, pp. 12-21. Springer.
547 Tachibana, R.O., Oosugi, N. & Okanoya, K. (2014) Semi-automatic classification of birdsong elements
548   using a linear support vector machine. *PloS one,* **9,** e92584.
549 Tinbergen, N. (1960) Comparative Studies of the Behaviour of Gulls (Laridae): a Progress Report1.
550   *Behaviour,* **15,** 1-69.
551 Walker, J.S., Jones, M.W., Laramee, R.S., Bidder, O.R., Williams, H.J., Scott, R., Shepard, E.L. & Wilson,
552   R.P. (2015a) TimeClassifier: a visual analytic system for the classification of multi-dimensional
553   time series data. *The Visual Computer,* **31,** 1067-1078.
554 Walker, J.S., Jones, M.W., Laramee, R.S., Holton, M.D., Shepard, E.L., Williams, H.J., Scantlebury, D.M.,
555   Nikki, J.M., Magowan, E.A. & Maguire, I.E. (2015b) Prying into the intimate secrets of animal
556   lives; software beyond hardware for comprehensive annotation in 'Daily Diary'tags.
557   *Movement ecology,* **3,** 1.
558 Williams, H.J., Holton, M.D., Shepard, E.L., Largey, N., Norman, B., Ryan, P.G., Duriez, O., Scantlebury,
559   M., Quintana, F. & Magowan, E.A. (2017) Identification of animal movement patterns using tri-
560   axial magnetometry. *Movement ecology,* **5,** 6.
561 Wilson, R., Culik, B., Bannasch, R. & Lage, J. (1994) Monitoring Antarctic environmental variables using
562   penguins. *Marine Ecology Progress Series,* **106,** 199-202.
563 Wilson, R.P., Shepard, E. & Liebsch, N. (2008) Prying into the intimate details of animal lives: use of a
564   daily diary on animals. *Endangered Species Research,* **4,** 123-137.

565 Ydesen, K.S., Wisniewska, D.M., Hansen, J.D., Beedholm, K., Johnson, M. & Madsen, P.T. (2014) What
566   a jerk: prey engulfment revealed by high-rate, super-cranial accelerometry on a harbour seal
567   (Phoca vitulina). *Journal of Experimental Biology,* **217,** 2239-2243.
568 Yoda, K., Kohno, H. & Naito, Y. (2004) Development of flight performance in the brown booby.
569   *Proceedings of the Royal Society of London B: Biological Sciences,* **271,** S240-S242.
570 Yoda, K., Sato, K., Niizuma, Y., Kurita, M., Bost, C., Le Maho, Y. & Naito, Y. (1999) Precise monitoring of
571   porpoising behaviour of Adélie penguins determined using acceleration data loggers. *Journal*
572   *of Experimental Biology,* **202,** 3121-3126.
573

574

Supplementary Material 1: LoCoD method operation

Each command the user selects has an opcode and possibly an operand. As the user selects various commands, a list of these opcodes/operands is stored in order of entry. This is then parsed into Reverse-Polish-Notation (AKA Post-Fix):
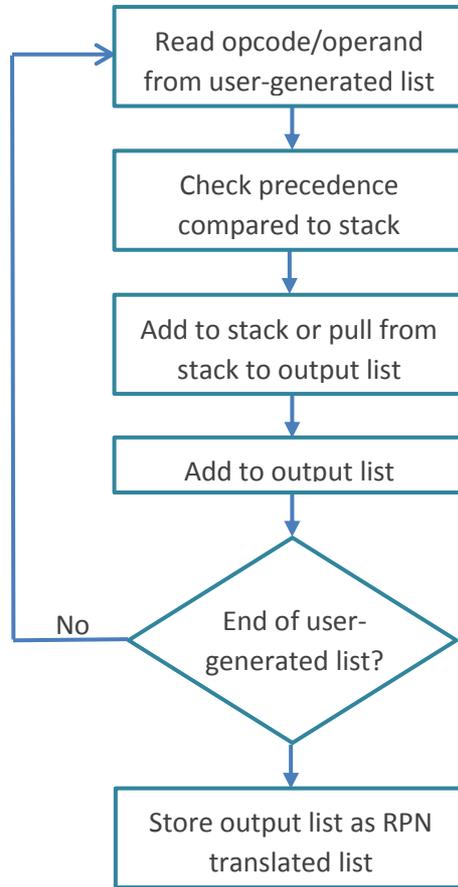
For example:

    If (Accel X Smoothed < 1.5) then "Mark-event"

    1,3,7,8,11,1.5,15,22

    1: If (

    3: Channel (7) "Accel X Smoothed"

    8: <

    11: Value (1.5)

    15:)

    22: Mark-event

This is simply translated to Post-fix as:

Accel X Smoothed (value of), Value 1.5, <

3,7,11,1.5,8

i.e. Post-fix would process this by reading the value of Accel X Smoothed and pushing this onto a stack, pushing value of 1.5 onto the same stack. The "<" command requires the two previously stored values on the stack. A "<" comparison between these two values results in a True or False.

622
623 The Time-Series algorithm:
624 Definition: *ETNE* is '*Extend To Next Event*', which is where an *Element's* validity is checked beyond
625 its stated valid range. It is checked from the starting *Event* number to the end of *Event* number + *Range*
626 + *Flex* i.e. as far as the next *Element* might exist. The point where it fails (if at all) is stored
627 For every *Event* stored in memory, each *Element* is parsed and the result stored
628 Once all Elements have been parsed through all in-memory data, the program checks if the TS
629 expression passes for each *Event*

- For all *n Elements*, beginning from the first *Event*, the program begins with *Element* 1
- The program checks if there are *Element$_i$* (*Valid*) consecutive points beginning at *Event n* to *n+Valid*
- If *Element$_i$* has passed, the program then checks if *Element$_i$* has *ETNE* enabled; if so, the program also checks from *Event n* to *Event n+Range+Flex* and stores the point of fail, or simply *n+Range+Flex* if no fail occurred
- If parsing *Element* > 1, the program then checks if the previous *Element* had *ETNE* as part of its construct. If so, it checks at which datapoint the previous *Element* failed. If it failed before the point the current *Element* passed, then the current *Element* fails.
- If the current *Element* failed, the program then moves onto *Event n+1* and starts again with *Element* 1
- If the current *Element* passed, the program then moves onto the next *Element*
- If all *Elements* have passed, the program then *Bookmarks* from the first *Element's Event* to the last *Element's Event + Valid* width; it then moves point *n* onto the end of the *Bookmark* just created as this behaviour's existence has been confirmed.

645
646

647 **Supplementary Material 2: Behaviour description in terms of LoCoD**
648
649 Any behaviour can be described by the sequence of defined motions, each motion defining a base
650 element of the behaviour. Each base element differs in the time over which it is performed and hence so
651 does the entire duration of the behaviour. The examples shown here have been selected as they differ in
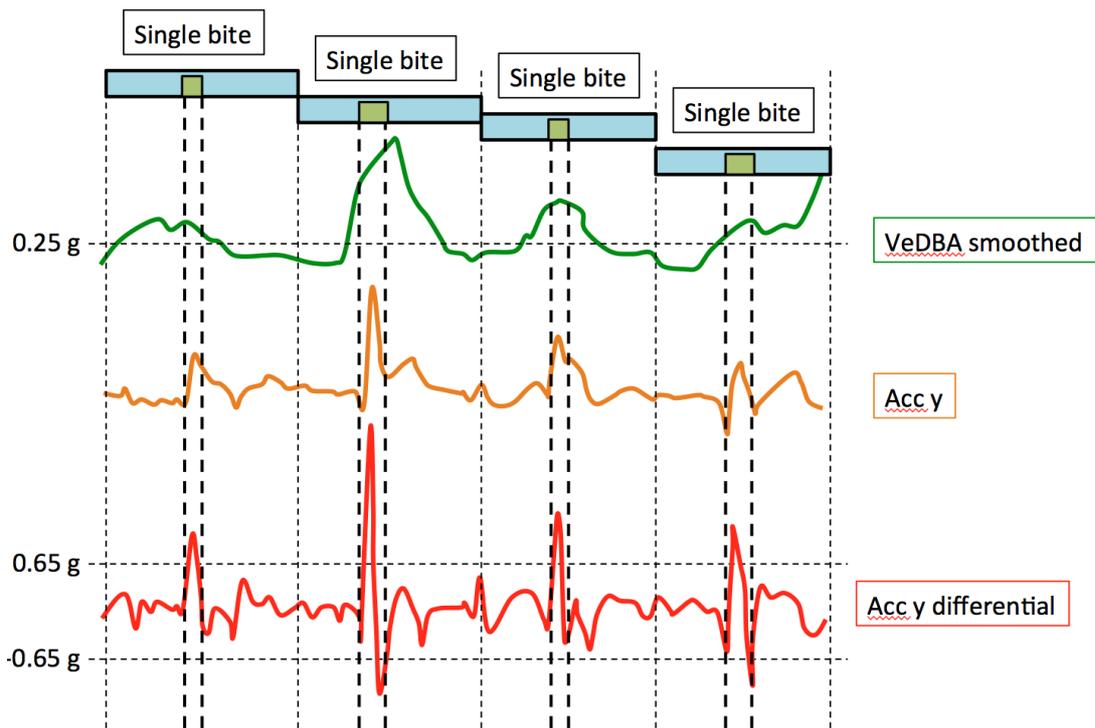652 the type and number of base elements involved as well as the duration of the behaviour.
653
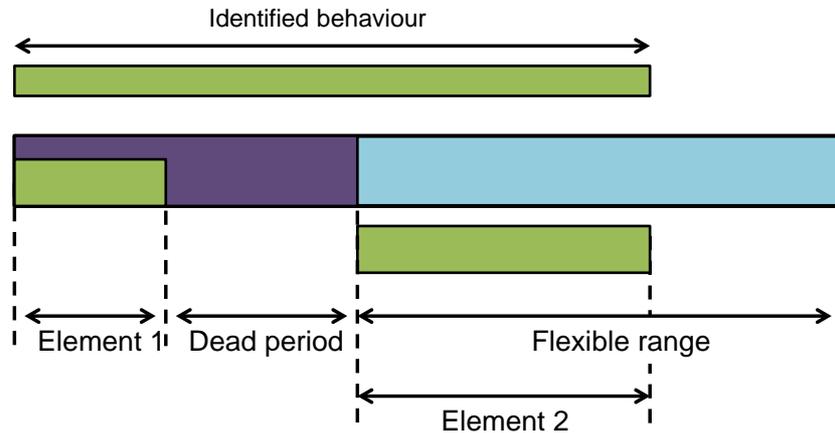
654
655
656 **FIGURE 1** Schematic diagram of a behavior in terms of; behavioural elements, a 'dead period' (see
657 text) and a flexible range of time within which behavioural elements must follow for all the behavioural
658 elements to constitute one LoCoD
659
660 **Short-period behaviours**
661
662 *Sheep biting*
663 For sheep and most herbivores, grazing is a complex behaviour that can be decomposed into smaller
664 behaviours such as biting and chewing. Biting consists in the extraction of the foodstuff from the
665 environment and chewing is the first stage of food processing. Herbivore bites are typically short and
666 high frequency behaviours that can occur in periods of less than a second. Biting is commonly
667 characterized by an abrupt head movement that typically occurs in one of two main directions; forward
668 and backward, which is also accompanied by an increase in standard movement metrics (e.g. VeDBA).
669 These abrupt head movements indicative of biting are well represented in the surge axis of the
670 acceleration, and the differential of this signal can be smoothed to reduce the influence of overall
671 motion of the animal while feeding. Although sheep biting is a short-period behaviour, the duration and
672 frequency of bites can be variable according the type of vegetation that individuals consume. For this
673 reason, we included a flexibility window of 10 consecutive data points (corresponding to 0.25 s if the
674 data are recorded at 40 Hz) to be able to detect this variability (see supplementary information 2 for
675 detail).

Single bite

Single bite

Single bite

Single bite

0.25 g

VeDBA smoothed

Acc y

0.65 g

Acc y differential

-0.65 g

FIGURE 2 Schematic diagram to demonstrate how the single bites of a sheep can be defined within the BE and flexible search criteria (colour coding for these as in Fig. 1). For precise details, see supplementary information 2. Four single bites are shown here as performed in sequence.

Walking Penguin

In contrast to the dive, the signal created by a penguin as it walks is comparatively short-period and complex, yet highly stylised in its pattern of motion. As the penguin makes a double step in walking, it sways from side to side, creating peaks and troughs in the smoothed signal in the sway axis of the acceleration. The differential of this signal can be smoothed again to reduce the noise manifest in effects of style or gait on the overall motion of the behaviour and can thus be used to classify all examples of a double step in walking. Differences in speed will still be apparent however, and so the use of a time flexible algorithm to classify the behaviour is ideal (see supplementary information 2 for detail).
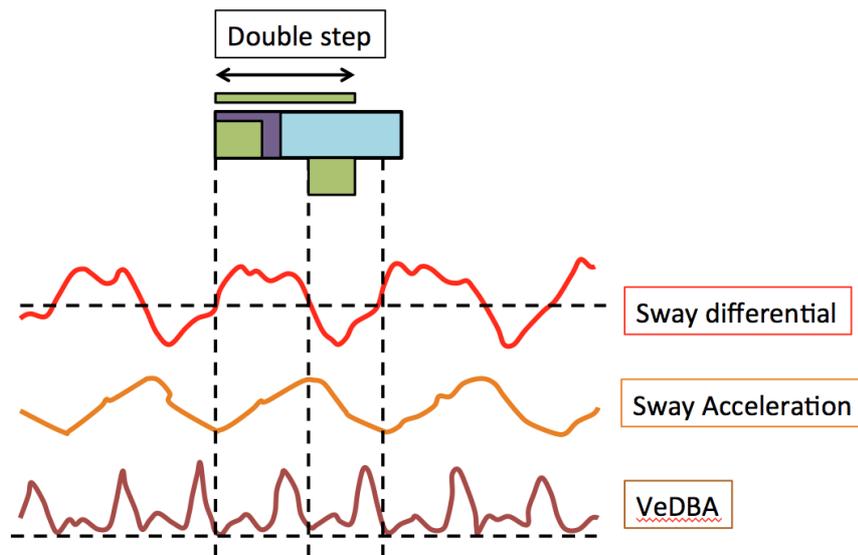
FIGURE 3 Schematic diagram to demonstrate how walking by a penguin can be defined within various
BEs, dead elements and flexible search criteria (colour coding for these as in Fig. 1). For precise details,
see supplementary information 2.

Working example with penguin walking

We present a step-by step example of the application of the LoCoD method to label walking behaviour
within a section of data recorded from the device attached to the Magellanic Penguin (penguin
walking_data section.raw). This includes a video attached (penguin walking.mp4) of the precise actions
undertaken during the process. [Note that the program provided can load acceleration data even if they
are not derived from the 'Daily Diary' provided that the data are arranged in columns with TAB as a
separator. In this case, the filename must be Xxx.col and under 'file of type', the 'col' section needs
using. In this case, you will be asked to specify sampling rate.]
Firstly, walking is identified as the signal in Figure 3 and a pattern of change identified in the smoothed
y-acceleration. The specific process (mirrored in the video) is;

1. Load raw data file with 40 Hz sampling frequency
2. Smooth the acceleration
3. Derive the differential of the smoothed y-acceleration
4. Use the 'display overlay window' to establish thresholds from templates
5. Define the base element equations based on the values in the display
6. Open 'build a time series based behaviour' and define time series windows
7. Bookmark matches to the template
8. Further rules can be applied to improve classification accuracy. In this case, walking can be
   regarded as a continuous behaviour so we merge bookmarks that occur within 80 data points (2
   seconds at 40 Hz sampling) and then remove bookmarks that are fewer than 80 data points in
   duration.
9. Note that all walking is correctly identified by this process.
10. These bookmarks can be exported as a master txt file for analyses in other software.

723   **Medium-period behaviour**

724

725   *Themalling Condor*

726   During thermal soaring, a condor must make a series of complete rotations to maintain a position within

727   the thermal and rise in the updraft. Each complete rotation can easily be seen in the magnetometer data

728   as the bird turns through all headings in relation to magnetic north. Hence each complete turn is defined

729   by a sine wave pattern in the x-axis of the magnetometer sensor, the length of which depends on the

730   time it takes for the bird to complete the turn. The behaviour is also expected to increase in duration

731   from seconds to minutes through a single climb and with increasing thermal strength and so this

732   behaviour lends itself to classification with temporal flexibility rather than any restricted classification

733   by pair-wise correlation, for example. In terms of classification the sine wave can be reduced to two

734   base elements, the first and second halves of the complete turn (see supplementary information 2).
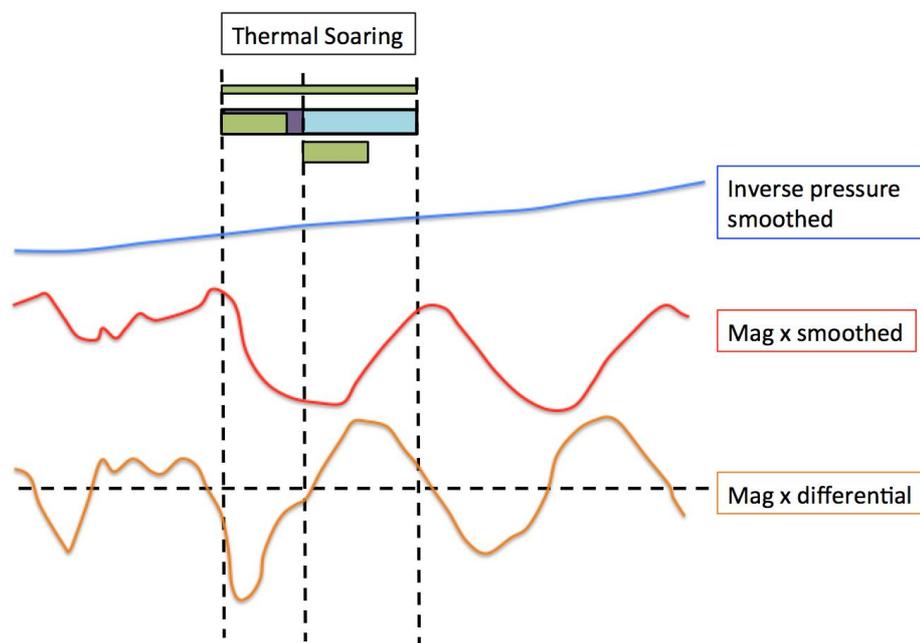
735

736   **FIGURE 4** Schematic diagram to demonstrate how thermal soaring by a condor can be defined within

737   various BEs, dead elements and flexible search criteria (colour coding for these as in Fig. 1) using

738   patterns in the output from the magnetometer and barometric pressure sensor. For precise details, see

739   supplementary information 2.

740

741

Long-period behaviour

Cheetah Stalking

744 When a cheetah stalks its prey, it reduces the acceleration signal in its movement, crouching low to the
745 ground, moving slowly closer to its prey. Thus, in terms of signal outputs, the rate of change of
746 smoothed acceleration defines the stalk poorly as there is very little change in the animal's postural
747 orientation. Instead, the defining feature is a lack of variation in any of the three acceleration signals
748 and hence a consistently low VeDBA. In this stalking phase, as the animal moves in on its prey,
749 changes in the smoothed magnetometry signals may also be evident although the rate of change in
750 directional orientation is not specific to the behavior. The chase follows the low VeDBA stalk
751 immediately. This is characterized by sprinting and a dramatic increase in the dynamism of movement,
752 resulting in an extremely high VeDBA relative to other behaviours. Stalking behavior in the cheetah
753 can therefore be identified using the two BEs that make up the LoCoD; i) a low VeDBA stalk, followed
754 by ii) a high VeDBA chase, each BE lasting at least several seconds.

755
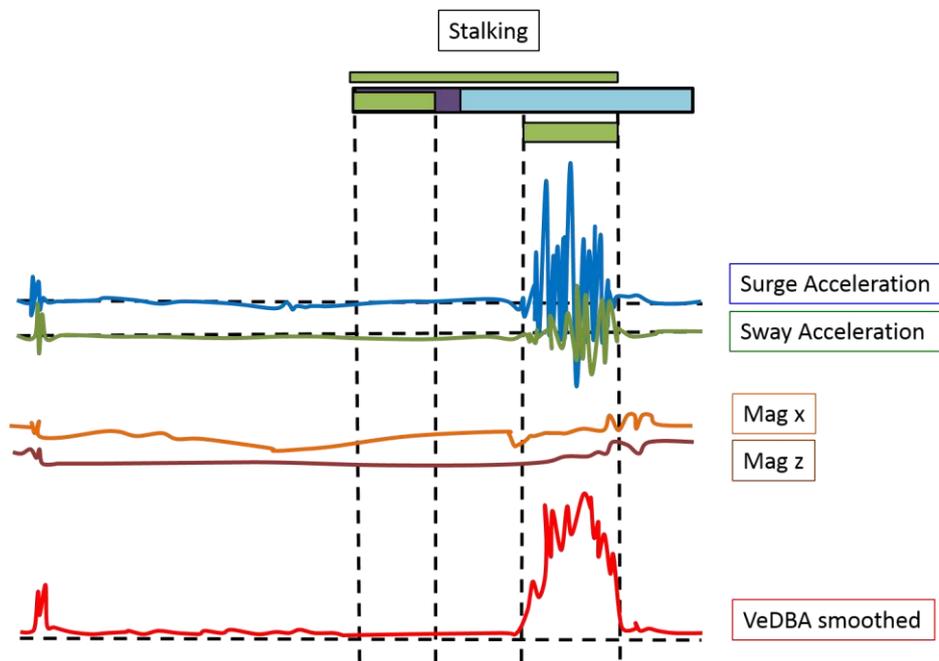
756
757
758 **FIGURE 5** Schematic diagram to demonstrate how stalking by a cheetah can be defined within various
759 BEs, dead elements and flexible search criteria (colour coding for these as in Fig. 1). For precise details,
760 see supplementary information 2.

761
762

763 **Supplementary Material 3: LoCoD method algorithm design**

764 Table S3.1: LoCoD method algorithm design for sheep biting. The different design components showed
765 in this table are; the variables used for processing, the base elements identified, and the time series of
766 those base elements. Numeric values shown refer to numbers of consecutive data points recording at 40
767 Hz so that, for example, the smoothing window is over 1 s.

Sheep Biting – short-period behaviour				
Processing	**Signal**	**Smoothing window**	**Differential range**	
	VeDBA	40	-	
	Acc y	2	5	
Base elements	Bite	If (VeDBA smoothed > 0.25) AND ABS(Diff_Accel y)>0.65) then mark events *(Include forward and backward head movement by using* ABS()		
Time series	**Element**	**Present**	**range**	**Flexibility**
	1 Head movement (Forward or Backward)	1	-	10

768
769

770 Table S3.2: LoCoD method algorithm design for penguin walking. The different design components
771 showed in this table are; the variables used for processing, the base elements identified, and the time
772 series of those base elements. Numeric values shown refer to numbers of consecutive data points
773 recording at 40 Hz.

Penguin Walking – short-period behaviour				
Processing	**Signal**	**Smoothing window**	**Differential range**	
	Acc y	10	5	
Base elements	Step left	If (SM (Diff_Accel Y smooth, 5) < -0.1) then mark events		
	Step right	If (SM (Diff_Accel Y smooth, 5) > 0.1) then mark events		
Time series	**element**	**present**	**range**	**Flexibility**
	1 Step left	6	16	16
	2 Step right	6	-	-

774

775

776 Table S3.3: LoCoD method algorithm design for condor thermalling. The different design components
777 showed in this table are; the variables used for processing, the base elements identified, and the time
778 series of those base elements. Numeric values shown refer to numbers of consecutive data points
779 recording at 40 Hz.

Condor Thermalling – medium-period behaviour				
Processing	**signal**	**Smoothing window**	**Differential range**	
	pressure	830	200	
	Mag x	40	80	
Base elements	½ turn section 1	if((smooth(diff_mag_x_smooth,20)>0) AND (diff_pressure_smooth<0))then mark events		
	½ turn section 2	if((smooth(diff_mag_x_smooth,20)<0) AND (diff_pressure_smooth<0))then mark events		
Time series	**element**	**present**	**range**	**Flexibility**
	1 ½ turn section 1	200	400	200
	2 ½ turn section 2	200	-	-

780
781

782 Table S3.4: LoCoD method algorithm design for cheetah stalking. The different design components
783 showed in this table are; the variables used for processing, the base elements identified, and the time
784 series of those base elements. Numeric values shown refer to numbers of consecutive data points
785 recording at 40 Hz.

Cheetah Stalking – long-period behaviour				
Processing	**signal**	**Smoothing window**	**Differential range**	
	VeDBA	10	NA	
Base elements	Stalk	If (SM (VeDBA Smoothed, 5) < 0.5) then mark events		
	Chase	If (SM (VeDBA Smoothed, 5) > 0.55) then mark events		
Time series	**element**	**present**	**range**	**Flexibility**
	1 Stalk	400	600	1200
	2 Chase	340	-	-

786

787 **Supplementary Material 4: LoCoD and Machine learning performance**

788 Here, we provide a brief description of each machine learning algorithm available in AccelerRater:

789 K-Nearest neighbors: This is a non-parametric method that labels a new sample/observation using a
790 vote between the K points in the training data set nearest to it. The method is a primitive form of
791 machine learning that is often referred to as 'lazy learning' because induction occurs during run time.
792 By default, we set K=3. For more detail, see James et al. (2013) and Bidder et al. (2014).

793 Linear SVM: Linear support vector machines compute the maximum margin hyperplane between two
794 classes. The multi-class extension used computes such a hyperplane between every two classes and uses
795 a vote to determine the class for a new point quantifying the similarity of a pair of observations using
796 Pearson correlation. More detail is provided by James et al. (2013).

797 RBF kernel SVM: This model is similar to a Linear SVM, but instead of using Gaussian kernels
798 employs Radial Basis Functions (RBF) kernels. The algorithm automatically determines centres,
799 weights and thresholds that minimize an upper bound on the expected test error. See Scholkopf et al.
800 (1996) for more detail.

801 Decision tree: This is a probabilistic method that works on binary decisions that are constructed
802 hierarchically. Basically, this method consists of a set of hierarchical decision rules developed to predict
803 the class of unclassified samples. Each rule can branch into another rule or a terminal category.

804 Random forest: This method consists of a combination of decision trees where each classifier is
805 generated using a random vector sampled independently from the input vector. This means that the
806 procedure is similar to a decision tree but includes introduced stochasticity. Instead of potentially using
807 all the variables to determine the best split at each node, only a randomly selected subset of variables is
808 used. For more detail, see Breiman (1999) and Breiman (2001).

809 Naïve Bayes: The Naïve Bayes algorithm is a simple probabilistic classifier that calculates a set of
810 probabilities by counting the frequency and combinations of values in a given data set. The algorithm
811 uses Bayes theorem and has a strong assumption that all attributes are independent given the value of
812 the class variable (i.e., features are conditionally independent). More detail is given in Patil & Sherekar
813 (2013).

814   LDA: The Linear Discriminant Analysis method is basically a linear model assuming Gaussian
815   distributions with equal covariance. See James et al. (2013) for more detail.

816   QDA: The Quadratic Discriminant Analysis method is the same as LDA, but without assuming equal
817   covariance (i.e., assumes that each class has its own covariance matrix). For more information, see
818   James et al. (2013).

819   ANN: Artificial Neural Networks (ANNs) are computer-based algorithms that imitate the structure and
820   behavior of neurons in the human brain. These algorithms can be trained to recognize and categorize
821   complex patterns. Pattern recognition is achieved by adjusting parameters of the ANN by a process of
822   error minimization through learning from experience. They can be calibrated using any type of input
823   data and the output can be grouped into any given number of categories. More detail is given in Bishop
824   (1995).

825

826   References:

827   Bidder, O. R., Campbell, H. A., Gómez-Laich, A., Urgé, P., Walker, J., Cai, Y., ... & Wilson, R. P.
828   (2014). Love thy neighbour: automatic animal behavioural classification of acceleration data using the
829   k-nearest neighbour algorithm. PloS one, 9(2), e88609

830   Bishop, C.M. 1995. Neural Networks for Pattern Recognition. Clarendon Press, Oxford.

831   Breiman, L. 1999. Random forests - Random Features. Technical Report 567, Statistics Department,
832   University of California, Berkeley, ftp://ftp.stat.berkeley.edu/pub/users/breiman.

833   Breiman, L. 2001. Random forests. Machine learning, 45(1), 5-32.

834   Scholkopf, B., Sung, K. K., Burges, C. J., Girosi, F., Niyogi, P., Poggio, T., & Vapnik, V. 1997.
835   Comparing support vector machines with Gaussian kernels to radial basis function classifiers. IEEE
836   transactions on Signal Processing, 45(11), 2758-2765.

837   James, G., Witten, D., Hastie, T., & Tibshirani, R. 2013. An introduction to statistical learning. New
838   York: springer.

839   Patil, T. R., & Sherekar, S. S. 2013. Performance analysis of Naive Bayes and J48 classification
840   algorithm for data classification. International Journal of Computer Science and Applications, 6(2),
841   256-261.

842

843

844 Table S4.1: LoCoD and Machine learning performance for sheep biting. Where performance is
845 measured in terms of the number of True Positive (TP), True Negative (TN), False positive (FP) and
846 False Negative (FN) results and the performance metrics of Recall (R), Precision (P) and Accuracy (A)
847 have been calculated. Note that for the Machine learning methods, each data point is labelled, so it is
848 possible to assign a category of TN. However, for the LoCoD method, data is labelled within the
849 LoCoD, so it is not possible to assign a category of TN as a non-existent LoCoD cannot be falsely
850 identified. For the latter, an accuracy value cannot be calculated.

behaviour	**Sheep biting**						
	Time (s)	Performance			Cases		
		R	P	A	TP	FN	FP
Manual	2039	1	1	1	171	0	0
LoCoD	**1.5**	**0.887**	**0.871**	**NA**	156	35	29
Nearest Neighbour	243	0.000	0.000	0.998	0	0	171
Linear SVM	3189	0.000	0.000	0.998	0	0	171
RBF SVM	253	0.000	0.000	0.998	0	0	171
Decision Tree	242	0.000	0.000	0.997	0	171	0
Random Forest	281	0.000	0.000	0.998	0	171	0
Naïve Bayes	317	0.000	0.000	0.998	0	0.0171	171
LDA	264	0.000	0.000	0.998	0	0	171
QDA	353	0.988	0.002	0.977	169	3	75
ANN	3451	0.000	0.000	0.988	0	0	171

851

852

853 Table S4.2: LoCoD and Machine learning performance for penguin walking represented by single
854 steps. Where performance is measured in terms of the number of True Positive (TP), True Negative
855 (TN), False positive (FP) and False Negative (FN) results and the performance metrics of Recall (R),
856 Precision (P) and Accuracy (A) have been calculated. Note that for the Machine learning methods, each
857 data point is labelled, so it is possible to assign a category of TN. However, for the LoCoD method,
858 data is labelled within the LoCoD, so it is not possible to assign a category of TN as a non-existent
859 LoCoD cannot be falsely identified. For the latter, an accuracy value cannot be calculated.

behaviour	Penguin walking						
	Time (s)	Performance			Cases		
		R	P	A	TP	FN	FP
Manual	2040	1.000	1.000	1.000	343	0	0
LoCoD	**14**	**0.982**	**0.984**	**NA**	335	8	8
Nearest Neighbour	77	0.971	0.965	0.973	337	6	6
Linear SVM	359	1.000	0.752	0.862	343	0	81
RBF SVM	79	0.939	0.973	0.964	322	21	6
Decision Tree	80	0.965	0.964	0.971	331	12	9
Random Forest	82	0.979	0.964	0.976	336	7	9
Naïve Bayes	75	0.992	0.761	0.866	340	3	0
LDA	74	0.988	0.762	0.867	343	0	78
QDA	77	0.759	0.709	0.771	261	82	76
ANN	405	0.925	0.966	0.947	317	26	13

860

861

862

863 Table S4.3: LoCoD and Machine learning performance for condor thermalling. Where performance is
864 measured in terms of the number of True Positive (TP), True Negative (TN), False positive (FP) and
865 False Negative (FN) results and the performance metrics of Recall (R), Precision (P) and Accuracy (A)
866 have been calculated. Note that for the Machine learning methods, each data point is labelled, so it is
867 possible to assign a category of TN. However, for the LoCoD method, data is labelled within the
868 LoCoD, so it is not possible to assign a category of TN as a non-existent LoCoD cannot be falsely
869 identified. For the latter, an accuracy value cannot be calculated. The machine learning methods
870 presented in this table are those that could be completed within 5 hours.

871

behaviour	Condor thermalling						
	Time (s)	Performance			Cases		
		R	P	A	TP	FN	FP
Manual	2220	1	1	1	146	0	0
LoCoD	**9**	**0.87**	**0.73**	**NA**	127	19	47
Nearest Neighbour	2182	0.144	0.257	0.797	21	125	11
Decision Tree	2358	0.006	0.355	0.838	1	145	0
Random Forest	2998	0.000	0.000	0.840	0	146	0

872

873

874 Table S4.4: LoCoD and Machine learning performance for cheetah stalking. Where performance is
875 measured in terms of the number of True Positive (TP), True Negative (TN), False positive (FP) and
876 False Negative (FN) results and the performance metrics of Recall (R), Precision (P) and Accuracy (A)
877 have been calculated. Note that for the Machine learning methods, each data point is labelled, so it is
878 possible to assign a category of TN. However, for the LoCoD method, data is labelled within the
879 LoCoD, so it is not possible to assign a category of TN as a non-existent LoCoD cannot be falsely
880 identified. For the latter, an accuracy value cannot be calculated. The machine learning methods
881 presented in this table are those that could be completed within 5 hours.

882

behaviour	*Cheetah stalking*						
	Time (s)	Performance			Cases		
		R	P	A	TP	FN	FP
Manual	180	1	1	1	10	0	0
LoCoD	**7.2**	**0.89**	**0.89**	**NA**	8	1	1
Nearest Neighbour	4045	0.996	0.986	0.983	10	0	9
Decision Tree	3470	0.999	0.986	0.985	10	0	10
Random Forest	4217	1	0.985	0.985	10	0	10
Naïve Bayes	3179	0.189	0.030	0.897	2	8	1
LDA	3016	0.056	0.259	0.984	9	0	9

883

884