OPEN

Quantifying free behaviour in an open field using k-motif approach

Marein Könings¹, Mark Blokpoel^{1,2}, Katarzyna Kapusta³, Tom Claassen¹, Jan K. Buitelaar³, Jeffrey C. Glennon³ & Natalia Z. Bielczyk^{3*}

Quantification and parametrisation of movement are widely used in animal behavioural paradigms. In particular, free movement in controlled conditions (e.g., open field paradigm) is used as a “proxy for indices of baseline and drug-induced behavioural changes. However, the analysis of this is often time- and labour-intensive and existing algorithms do not always classify the behaviour correctly. Here, we propose a new approach to quantify behaviour in an unconstrained environment: searching for frequent patterns (k-motifs) in the time series representing the position of the subject over time. Validation of this method was performed using subchronic quinpirole-induced changes in open field experiment behaviours in rodents. Analysis of this data was performed using k-motifs as features to better classify subjects into experimental groups on the basis of behaviour in the open field. Our classifier using k-motifs gives as high as 94% accuracy in classifying repetitive behaviour versus controls which is a substantial improvement compared to currently available methods including using standard feature definitions (depending on the choice of feature set and classification strategy, accuracy up to 88%). Furthermore, visualisation of the movement/time patterns is highly predictive of these behaviours. By using machine learning, this can be applied to behavioural analysis across experimental paradigms.

Open Field Paradigms In Translational Psychiatry

Rodent behavioural research analyses patterns of behaviour as face and predictive validity measures for human psychiatric and neurological conditions. These behavioural paradigms are relevant to a number of cognitive disorders, including schizophrenia, depression, bipolar disorder¹, anxiety and autism². The open-field paradigm is commonly deployed for the study of movement, learning, sedation and anxiety-related aspects³. In the open-field paradigm, an animal is placed in a constrained environment (a well-lit circular, square or rectangular area, bounded either by insurmountable walls or deep gaps) which it is free to explore. Its behaviour is then recorded over time. The characteristics of the environment can differ between experiments and may deploy objects to study interaction between the animal and object to examine reactions to novel and stressful situations. An open field with objects can also be used to examine repetitive checking behaviour, an important aspect of obsessive-compulsive behaviours⁴. Commonly recorded variables include horizontal locomotion (based on a count of transitions between marked areas within the field, vertical activity (based on the animal's rearing and leaning behaviour), as well as the latency (time) spent in certain areas and more specific behaviours such as head shakes and grooming³. Automated systems record all of the necessary variables with high precision and flexibility⁵, by recording the animal's position using a video camera and processing the resulting footage with specialised software such as Noldus EthoVision 3.0⁶. The open field is also often divided into a grid of zones for the analysis which can be of differing sizes dependent on the size of the open field. Typically this is divided into 5 by 5 zones. A range of experimental designs is used within the open-field paradigm. Subjects are divided into experimental groups which undergo different interventions. Differences between groups in terms of observed open-field behaviour are then assumed to result from the difference in intervention deployed. These experimental interventions include changes to the objects in the open field, introduction of another animal, differences in the amount/type of food available or administration of drugs expected to induce behavioural changes⁷.

¹Radboud University Nijmegen, Comeniuslaan 4, 6525 HP, Nijmegen, The Netherlands. ²Donders Centre for Cognition, Montessorilaan 3, 6525 HR, Nijmegen, The Netherlands. ³Department of Cognitive Neuroscience, Radboud University Nijmegen Medical Centre, Geert Groteplein Zuid 10, 6525 GA, Nijmegen, The Netherlands. *email: natalia.bielczyk@gmail.com

Automated Methods For Quantifying Rodent Behaviour

To accurately assess the differences in behaviour between experimental groups, the objective quantification of the behaviour is important. Although currently available software such as Noldus Theme can extract raw behavioural variables from video footage, little has been done to automate and optimize further analysis of the collected datasets. These approaches detect *T-patterns*, which denote the repeating sequences of fields visited by the animal during the experiment^{8,9}. This method while useful also involves a large number of significance tests, thereby increasing the risk of false-positive results¹⁰. In addition, the T-patterns method assumes that all areas of the open field are equally likely to be visited. In reality, the opposite is true: some areas are visited far more often than others. Furthermore, zone transitions are computed using the T-pattern method with no regard to the temporal relation between crossings. That is, the time that a subject spends within a zone is of no consequence, only the order in which zones are visited, counts. This is problematic as this does not enable a distinction between behaviours that are time-sensitive (e.g. spending time near an object as opposed to just passing by it). This is a consequence of the T-pattern method performing only spatial and not temporal segmentation. Finally, the T-pattern method only utilises the absolute position of the subject and does not account for other factors such as objects in the open field or the distance between the animal and the open field boundary. As such, behaviours dependent on the relative location of these objects cannot be detected using T-patterns. If the location of an object changes, for example between subjects, the T-pattern method cannot discover a pattern if both subjects display the same behaviour relative to the object but in different absolute positions. Lastly, the T-pattern method is only available as a part of commercial software, *Theme* by Pattern Vision (<http://patternvision.com/products/theme/>), which is a black-box software with a basic interface and does not allow for accessing the source code.

The K-Motif Approach

Here we sought to design a new method for quantifying behaviour in the open environment. In an attempt to overcome the above issues, we proposed a method which takes into account the temporal dynamics of the movement as well. This method adapts the k-motifs algorithm to search for the hierarchical structure of repeating patterns in behaviour. The most common patterns were then used to train a classifier to distinguish between experimental groups between quinpirole and vehicle-treated rats in the open field. Here, we improve the classification of experimental groups (further referred to as classes) based on behavioural differences in the open field.

In section 2.1, we provide a detailed description of the experimental datasets used for developing the classifier. In section 2.2, we introduce the k-motifs approach. Then, in section 3, we report the results of classification with use of k-motifs as features and compare this classification performance with classifiers using standard features. Lastly, in section 4, we critically discuss the results and potential applications for this methodology. We made all the tools developed for this paper accessible in open source.

We dedicate this work to researchers looking for data-driven biomarkers of behaviour in the open field. These patterns can help in better understanding of the characteristics of disorders related to movements, such as Obsessive-Compulsive Disorder, Tourette Syndrome, Parkinson's Disease, Huntington's Disease or Autism Spectrum Disorders. The k-motif method proposed in this manuscript can also help to measure effects of pharmacological therapy with higher sensitivity as the performance in classification between experimental and control group achieved with our method is the highest among all the available methods.

Materials and Methods

The experimental datasets. The experimental datasets include open field recordings in rats, collected at the Radboud University Nijmegen Medical Centre under ethical approval number DEC-2012-281. The sample included 6 subjects in the experimental group and 6 controls. All subjects were male Sprague Dawley rats (Charles River, Germany). The open field setup includes an open square area of size 160×160 [cm], placed 60 [cm] from the ground. An additional virtual circuit 20 [cm] wide around the table is added to the open field size to record behaviour where the subjects' head or tail reaches out over the table. The resulting area of 200×200 [cm] is divided into 25×40 [cm] zones. The area contained four small fixed objects (cubes or cylinders) to encourage exploration in animals (Fig. 1A, B). The objects were distributed as in¹¹, and included two black and two white objects. The experimental sample included 6 subjects in an experimental intervention group and 6 controls. In total, data was recorded from subjects across 13 30-minute sessions in the open field, which amounts to a total of 156 recorded sessions. Injection and training sessions took place every day for 13 consecutive days. Before each open field session, the subjects within the experimental intervention group received an injection of quinpirole (0.5 mg/kg i.p.). The behaviour in the open field was recorded with an overhead video camera. The resulting footage was then processed by a video tracking system Ethovision 3.0 from Noldus Information Technology BV, the Netherlands (Wageningen, The Netherlands)⁶, which produced a multivariate time series for each 30-minute session, representing multiple variables (Supplementary Table 1, Supplementary Material 1: *Experimental datasets*). The time resolution of the video recordings was 0.040[s]. The set of available variables is represented here by V . The collection of time series resulting from processing video footage from one session in one subject, corresponds to a single d in D (Eq. 1). Examples of the visualisations of the movement are presented in Fig. 1. Additionally, an exemplary visualisation containing single datapoints plotted every 1.0 second (every 40 data frames) is presented in Supplementary Fig. 5. This visualization better shows the dynamics as it better indicates that the subject spends most of the time around the objects placed in the open field. The datasets were preprocessed according to a preprocessing scheme is introduced in Supplementary Material 2: *Data preparation*. After cleaning the data, 10 out of 13 sessions, for 5 quinpirole subjects versus 3 controls were retained in the data (Supplementary Table 2), yet we augmented the datasets by sectioning every session into 10-minute segments (Supplementary Material 3).

$$d \in D = \{d_v | v \in V\} \quad (1)$$

Figure 1. Movement of example subjects during single sessions. (A) The spatial, ‘top-down’ view of the path taken by an exemplary control subject during an exemplary session. The borders of the open field are shown in red, with the four objects shown as red circles. (B) The same plotted for an exemplary subject from the quinpirole group. (C) The same path as shown in (A), x and y as in coordinates over time are shown as time series.

K-motif method. To address the limitations of T-patterns mentioned in Section 1, in this work, the k -motifs method is employed for feature extraction: finding recurring patterns^{12,13} in the data. This method involves the subsequent use of two methods: Symbolic Aggregate Approximation (SAX¹⁴) and Sequitur grammar induction algorithm¹⁵. The first part, SAX, segments the data both in the temporal domain and in the spatial domain. The second part, grammar induction using Sequitur, finds patterns in the resulting approximation by constructing a hierarchical grammar. The k -motifs algorithm successfully yields frequent patterns from a time series, which can be used to characterise the data. However, to enable classification of time series data, each series must be reduced to a feature vector. A feature vector is a list of numbers of a set length, where each number describes an aspect of the data. These numbers can then be compared between series, to determine their relative similarity. To use the patterns extracted by k -motifs, a conversion strategy to feature vectors must be devised. Additionally, while patterns in raw movement data are useful, even more patterns may be found by exploiting the animal’s position relative to different aspects of the open field. Finally, k -motifs produces a very large number of patterns, so to ensure the conciseness of the quantification method, some subset of the patterns must be selected for the final feature vector. All these topics will be discussed in the following sections.

Symbolic aggregate approximation (SAX). Discretizes the real-valued time series $x(t)$ of the animal x position (e.g., as seen in Fig. 1, blue series), starting by slicing it into segments of duration w . For each segment, the values contained within are averaged. This leads to a new time series $t(t)$ of length $\frac{|x|}{w}$. Note that if the length of $x(t)$ is not divisible by w (length of $x(t)$ is not a multiple of w), x is padded with the last value of x until its length is divisible by w (this is the case in the version of SAX used in our methods; see also generalization of the SAX code by^{16,17} which handles also lengths non-divisible by w).

Next, the range of the values of $t(t)$ is divided into many sections (further referred to as ‘bins’). The bins are defined using z-scores from a normal distribution with mean and standard deviation derived from the original data $x(t)$. We denote the number of bins by α , and each bin is represented by the lower limit in its value range. For example, if the values fall between -4.0 and 4.0 , and $\alpha = 4$, the bins are: $[-4.0, -2.0]$ represented by -4.0 , $[-2.0, 0.0]$ represented by -2.0 , $[0.0, 2.0]$ represented by 0.0 , and $[2.0, 4.0]$ represented by 2.0 .

Additionally, the bins are given unique labels. In this work, bins are labeled from 0 to $\alpha - 1$, using the ordering of the bins from the lowest to the highest. Then, each value in $t(t)$ is replaced by the label of the associated bin to

Figure 2. The SAX representation¹⁴. (A) The real-valued input data is presented in black. The numbers given on X-axis represent the subsequent data points, in which the real-valued input variable is measured. If the time resolution of the recording is high, this may be a near-continuous sequence. In this example, parameter w is set to such a value that the input is divided into segments of 1.2 time units. The mean value of the time series within each segment is used to produce the approximation sequence (presented in green). Finally, for each mean value, the corresponding bin is determined. In this example, there are four bins ($\alpha = 4$) delineated using dotted red lines. Bins width is determined in terms of standard deviations from the mean value. The final output in this example, is a sequence of labels 2 0 3 1 2. (B) Standard division into 25 zones (as implemented in the T-pattern method). (C) The division of open field with use of SAX ($\alpha = 10$, which means 10 bins along each dimension), with all positional data. Note the underlying normal distribution of visits is reflected in the sizing of the bins. In less popular areas, larger bins are available and less detail is recorded.

which belongs - namely, the highest bin whose lowest value is lower than the given value in $t(t)$. The obtained result is an approximate symbolic representation of the original time series. This transformation is formalised in Eq. 2. An example segmentation is demonstrated in Fig. 2A.

$$\text{SAX}(x, w, \text{bins}) = \left[i \text{ for } p \text{ in } [0, 1, \dots, \frac{|x|}{w} - 1] \text{ such that } \text{bins}_i \leq \sum_{pw \leq j < (p+1)w} \frac{x_j}{w} < \text{bins}_{i+1} \right] \quad (2)$$

where $i = \min(\text{bins}_i)$.

Assuming normal distribution of the input datasets, this ensures that each value in t is equally likely to fall into each bin, and thus that each event (falling into a bin 0 to $\alpha - 1$) occurs equally often. This is explained in Fig. 2, where $\alpha = 4$. In this example, the value range in the time series, is dissected into four equiprobable bins, centred on the mean of the time series. Note that, for any considered variable (e.g., the animal's x position on the open field table), the bins are precomputed using *all* available datasets: the bins are defined for the merged datasets from all the subjects within the experimental group. The purpose of this step is to have one, uniform output sentence for all subjects. An example set of bins, computed for all two-dimensional positional data in our datasets, is presented in Fig. 2B. An advantage of using bins based on the normal distribution derived from the datasets - rather than equally-sized bins - is that the proportion of visits in all bins becomes more uniform (this

s		$R(s)$	expanded	$c(s)$
S	$\rightarrow$	1 1	$a b c a b d a b c a b d$	1
1	$\rightarrow$	2 c 2 d	$a b c a b d$	2
2	$\rightarrow$	ab	$a b$	4

Table 1. Example grammar produced by Sequitur for the input sequence $a b c a b d a b c a b d$, demonstrating rules, expanded sequences and occurrence counts. Note that a, b, c, d are terminal symbols, and $S, 1$ and 2 are non-terminal symbols in this example.

may be evaluated by comparing the entropy of the visit counts, with a higher entropy meaning a more uniform distribution).

To apply the SAX algorithm, we reimplemented the original version of the algorithm using Python, the code is included in the open GitHub repository, at <https://github.com/MareinK/kmotifs-paper-code>.

Grammar induction using sequitur. The output of SAX algorithm can be interpreted as a series of symbols. Sequitur is an algorithm for inducing a hierarchical grammar from a given sequence of symbols, thereby compressing the input to a smaller representation¹⁵. This outcome small representation is, again, a series of symbols. Some symbols are *terminal*. A terminal symbol is a symbol that was also present in the original sequence. Other symbols are *non-terminal*. A non-terminal symbol is an outcome of the compression, and there is an associated *rule* which determines how to ‘expand’ this symbol to again a series of (non)-terminal symbols. A grammar is defined as a set of rules for repeatedly expanding all non-terminal symbols in this way will eventually reproduce the original sequence, with no more non-terminal symbols present. A special ‘start symbol’ (denoted by S in this text) determines the initial rule to use when starting the expansion. The Sequitur algorithm constructs a grammar from the input sequence by processing input symbols sequentially, and adding them to the grammar one by one. An example of an output grammar is given in Table 1. The symbol-to-rule mapping is given by $R:symbol \rightarrow sequence$.

The Sequitur algorithm processes input symbols sequentially, adding them to the grammar one by one, while maintaining two grammar properties:

- **Digram uniqueness:** Every pair of symbols (terminal or non-terminal) occurs *not more than once* in the grammar. If the same pair of symbols occurs twice at some point, a new rule is created mapping a non-terminal symbol to this sequence of two symbols, and the two occurrences of the pair are replaced by the new non-terminal symbol. This property ensures that the grammar gets compressed once a repetition is found. It also imposes a hierarchical structure on the grammar.
- **Rule utility:** Every rule is used *at least twice*. Because of the digram uniqueness property, a pair consisting of a non-terminal symbol and another symbol may be condensed as a new rule. This causes that non-terminal symbol's occurrence count drops by one, which might become the only occurrence of the non-terminal symbol. In this case, the rule is removed, and its non-terminal symbol is replaced by the corresponding sequence. This property not only ensures that no useless rules (with only one occurrence) exist, but also allows for formation of rules longer than two symbols, which would otherwise not be possible.

By maintaining these properties, and using the proper data structures, the Sequitur algorithm runs in linear time in the length of the input sequence¹⁵. The output grammar can then be used to determine the relative frequency of sub-sequences by their occurrence count, so that they may be used in the quantification method.

In the context of finding common patterns in the SAX output, we are interested in the *occurrence count* of the different non-terminal symbols, since these represent common sequences. Given a grammar, let s denote a non-terminal symbol in that grammar. Then we can find the occurrence count $c(s)$ of that symbol by finding all the rules that s occurs in, and taking the total occurrence count of the symbols associated with all those rules. S always has an occurrence of 1 (Eq. 3).

$$c(s) = \begin{cases} 1 & \text{if } s = S \\ \sum_{t \mid s \in R(t)} c(t) & \text{otherwise} \end{cases} \quad (3)$$

For example, in Table 1, the most frequent sub-sequence (longer than 1 symbol) is $a b$, as it occurs four times in the input. We introduced the occurrence count and added to the algorithm to apply Sequitur to this particular research problem, namely finding patterns in the SAX output data.

Feature vectors. The k -motifs method, using SAX and Sequitur, yields frequent motifs for a certain time series, which may be used to give insights into rodent behaviour. However, these motifs are not suited for classification since they do not take the form of a feature vector. Specifically, there is no obvious comparison between motifs derived from different time series.

To resolve this issue, the following method is applied: given a set of time series from which feature vectors need to be extracted, first find the k most interesting motifs (defined hereafter) for the complete dataset. We further refer to this set of interesting motifs as $\mathcal{M}_k(D)$ (Eq. 4). Now, to create a feature vector for a single time series, $\mathcal{F}_k(d \in D)$, we need to determine the *frequency* of each motif discovered in that time series (Eq. 5). This method yields feature vectors that are comparable for classification, element-wise.

$$\mathcal{M}_k(D) = k \text{ most interesting motifs in } D \quad (4)$$

$$\mathcal{F}_k(d \in D) = [\text{count of motif } m \text{ in } \mathcal{M}_k(D)] \quad (5)$$

In the context of this manuscript, the classification relates to discrimination between the experimental intervention group compared to the control (vehicle-treated) group (based on open field data). These groups, in machine learning terminology, are referred to as *classes*. In many classification problems, the uneven size of classes is an issue. To account for this, k frequent motifs should be extracted from the sessions of each class separately. Given n classes, this will yield $k \cdot n$ frequent motifs. As before, a feature vector representing $d \in D$ is created by determining the frequency of each of the motifs in the session. Thus, the length of the feature vector is dependent on the number of classes n as well as the number of extracted motifs k .

Interestingness. Using combinatorics, the number of possible motifs grows very fast with the length of the sequence. Therefore, one should define some rules for motif preference. This can be achieved by defining ‘interestingness’ of motifs. Given a motif m , three interesting properties may be defined, each of which one may wish to maximise when searching for interesting motifs. To evaluate a motif’s interestingness $I(m)$, some (non-linear) combination of these three properties can be used. The k most interesting motifs could then be selected as those with the highest interestingness.

- **Frequency $I_f(m)$:** The number of times this motif occurs in the data; the more frequent a motif is, the more suitable it will be as a feature. Of course, certain motifs which happen to be rare, can still carry meaningful information and be useful for classification but in practice, there is no efficient way to select these motifs because using rare motifs would facilitate over-fitting (i.e., building a classification tool that contains more parameters than can be justified by the data, and will fail to predict future observations).
- **Length $I_l(m) = |m|$:** the number of symbols in the motif. Longer motifs are more interesting because they expose more detailed patterns. A motif of length 2 only shows a linear movement from one position to another, while longer movements will describe more detailed aspects of the animal’s behaviour.
- **Diversity $I_d(m)$:** the number of unique symbols in the motif. A long motif may only contain a few unique symbols, which means the motif describes a movement with a high amount of repetition. Since such a pattern is clearly a combination of other motifs, it might be more interesting to focus on motifs which cannot be easily decomposed. This is encouraged by preferring motifs with a high number of unique symbols.

Certain relations hold between these properties. Given any motif m and symbol s , a new motif may be constructed as $n = m + s$ by appending s to m . Regardless of the choice of m and s it is clear that n cannot be more frequent than m , since for each occurrence of n there is also an occurrence of m contained within. Thus, $I_f(m + s) \leq I_f(m)$ and $|m + s| > |m|$ for any m and s . This means that frequency is inversely proportional to length: $I_f(m) \propto |m|^{-1}$. Note also that $|m| \geq I_d(m)$, meaning that length must grow with diversity, and so $I_f(m) \propto I_d(m)^{-1}$.

This means that, when using these three properties and k -motifs to define a quantification method, high accuracy (I_f) and high interpretability (I_l and I_d) are mutually exclusive. It is not clear whether this holds for quantification methods in general. In any case, the best way to combine the three properties into a single measure of interestingness $I(m)$ is not obvious. In this work, four different ways of combining the three measures are explored (Eqs. 6–9). While I_1 combines the three measures with equal weight, the other functions favour one measure while devaluing the other measures by scaling them logarithmically.

$$I_1(m) = I_f(m) \cdot |m| \cdot I_d(m) \quad (\text{equal weight}) \quad (6)$$

$$I_2(m) = I_f(m) \cdot \log(|m|) \cdot \log(I_d(m)) \quad (\text{focus on frequency}) \quad (7)$$

$$I_3(m) = |m| \cdot \log(I_f(m)) \cdot \log(I_d(m)) \quad (\text{focus on length}) \quad (8)$$

$$I_4(m) = I_d(m) \cdot \log(I_f(m)) \cdot \log(|m|) \quad (\text{focus on diversity}) \quad (9)$$

The choice of combination matters because they can result in different sets of leading motifs.

Spatial relations. Although the k -motifs algorithm may be applied to any time series, only the X centre and Y centre variables from V are used to construct the feature vectors. The choice to restrict the number of variables was made to reduce the size of the feature vector, thereby increasing interpretability. The choice for these specific variables was made since they are the variables used by other methods, e.g., T-patterns. Then, we should account not only for the subject’s position in space, but also for the relationship to *objects* and the *boundaries* of the open field. Therefore, the k -motifs algorithm must not only be applied to a time series tracking the subject’s position, but also to time series reparameterised concerning objects. Each of these transformed time series is called a *spatial relation* (including the original position).

This extension of the feature space further increases the length of the feature vector to $k \cdot n \cdot r$, with r the number of relations we wish to capture. The spatial relations that we chose, are presented in Table 2 (although of course, more possibilities exist). In our specific dataset, $n = 2$ (control versus quinpirole) and $r = 4$ (absolute

Name	Dimensions	Description
absolute position	2	2 location of the animal on the table in Cartesian coordinates
relative position	2	change in (x, y) location of the animal as compared to the previous time point
nearest object	1	the distance from the animal to the nearest object
nearest boundary	1	the distance from the animal to the nearest point on the boundary of the table

Table 2. Chosen spatial relations.

Component	Symbol	Domain	Description
k -motifs	w	$\mathbb{N}_{>0}$	segment size used for dividing the input sequence
k -motifs	α	$\mathbb{N}_{>0}$	number of bins
vector creation	k	$\mathbb{N}_{>0}$	number of most interesting motifs to select per class
vector creation	I	$\{I_1, I_2, I_3, I_4\}$	measure of interestingness

Table 3. An overview of the free parameters for the full k -motifs pipeline.

position, relative position, minimum distance to any object, minimum distance to the boundary). The length of the feature vector, then, becomes $8k$.

Note that some spatial relations are two-dimensional (absolute/relative position) while others are one-dimensional (distance). While using the k -motifs algorithm in application to one-dimensional data is straightforward, in case of two-dimensional data it becomes less obvious. We took the following approach. First, the two dimensions are processed by SAX individually, resulting in two separate symbolic time series. These two time series are then recombined to create a time series consisting of pairs of symbols. This new time series can then be processed by Sequitur, interpreting each pair as an individual symbol.

Free parameters. The complete quantification method, combining k -motifs, creation of the feature vector and spatial relations, involves several free parameters (Table 3).

Hill-climbing is a technique for optimising a cost-function characterising certain problem, by iteratively adjusting its input parameters¹⁸. This technique can be used to determine the optimal values in the parameter space (w , α and k), whereas exhaustive search is performed to optimise I . To evaluate the results of this optimisation, the complete dataset is split into two portions: 90% that is used for hill-climbing based on classifier performance (which in turn splits the data into training and test sets) and 10% that is held out for final evaluation (there are 10 sessions $\times$ 3 segments in each session = 30 data segments per animal which means that 3 segments per animal were saved for validation). To summarize, the schematic representation for the full k -motifs pipeline is given in Fig. 3.

Evaluation. Four popular classifiers are used to compare four different quantification methods. Implementation of these classifiers comes from sklearn package for Python¹⁹:

- **Gaussian Naive Bayes:** Assuming independence between features, Bayes' theorem can be applied to the data. This method learns probabilities with each combination of class and feature from the data, and from this evaluation, the most probable class can be inferred for a new observation²⁰.
- **Decision Tree:** In this method, several if-then-else decision rules are learned with the purpose to accurately divide the data into classes based on feature values. A new observation is then classified by evaluating the decision rules on its features²¹.
- **Multilayer Perceptron:** A feed-forward network of artificial neurons with nonlinear activation functions is used to model the relationship in the data between features and classes through connectivity weights between neurons. A new observation is then classified by feeding a set of features representing this observation into the network, which returns the predicted class at the output²².
- **k -Nearest Neighbours:** All observations in the data and their associated classes are stored in memory. Given a new observation, its class is determined by majority vote of the k data points closest to the observation in the feature space²³.

We chose the default parameters as implemented in the sklearn package. To determine the relative performance of classifiers, the label-frequency based macro F_1 score was used as a measure of test accuracy. As a variation of the commonly used F_1 score, this measure combines the classifier's precision P (also known as positive predictive value, PPV) and recall R (also known as sensitivity) with equal weights to evaluate the performance of the classifier. The F_1 score is computed individually for each class c , then the average between these scores is computed, with each score weighted by the size of the corresponding class $|c|$ (Eq. 10)²⁴. This ensures that imbalance in class size does not influence the score.

Figure 3. Schematic visualisation of the approach proposed in this work, based on k -motifs. The pipeline starts with the raw dataset (1), containing movement data from multiple open field sessions. Each of these sessions is then transformed into multiple spatial relations (2). For each spatial relation, the SAX bins are computed, using parameter α determining a total number of bins (3). The spatial relations are then separated based on the class of the session (4). For each class, every data point is processed using k -motifs, using the parameter w and the SAX bins computed earlier (5). For each data point, this results in a set of motifs, and all motifs of a class are gathered into a list. The motifs are sorted according to a measure of interestingness (I). Subsequently, certain number k of top motifs is selected to represent the class (6). The top motifs from each class are then combined to create the final set of motifs (7). These can then be used to create a feature vector for any data point (from the original data or otherwise), by counting the number of occurrences of each motif within that segment (8).

$$F_1^{\text{weighted}} = \sum_{c \in \mathcal{C}} \frac{2P(c)R(c)}{|c|(P(c) + R(c))} \quad (10)$$

To obtain reliable classification results, stratified k -fold cross-validation was used. In non-stratified k -fold cross-validation, the data is randomly divided into k subsets of equal size. Then, for each subsample i , the subsample is used as a test set while the other $k - 1$ subsamples are used for training. This results in k classifiers with each a classification score, which may be averaged to obtain the final score. This method ensures that all individual data samples are used for testing exactly once. Additionally, in *stratified* k -fold cross-validation, each subset is selected in a way that the distribution of classes is as near as possible to that of the complete data²⁵ (here, it means a ratio 1:1 between cases and controls). This prevents confounding the results by the class labels. In this project, $k = 10$ is used.

Comparison with other methods. In this work, we further compare k -motif approach with other methods which also define how a single experimental session is reduced to a simpler representation:

1. Established T-pattern method⁸ which interprets patterns as sequences of visited discrete fields (in this case, on the square experimental table divided into 25 square zones) within certain time-frame.
2. *Full data* method: no reduction is performed and the complete set of points visited during the session is taken as the representation of that session.
3. *Means and variances* method: the data of a session is reduced to just four numbers: the mean and the variance, per dimension (one of each for x-coordinates and y-coordinates).
4. *Behavioural* method: the data of a session is reduced by applying several hand-crafted heuristics to it that are thought to be useful descriptions of an animal's behaviour. This results in a collection of numbers that describe behaviour displayed during the session. For a full description of this method, check²⁶.

Ethical approval. The animal experiments in rats with use of quinpirole, were carried out following local guidelines and regulations within the Netherlands. All animal procedures were approved by the Ethical Committee on Animal Experimentation of Radboud University Nijmegen (under RU-DEC, number 2012-281) and the Dutch Ethical Committee on Animal Experimentation, and following the EU guidelines for animal experimentation.

Results

Normalising bins with SAX algorithm. Using the T-pattern method with linear zone divisions, the entropy of the proportion of visits per zone is 2.90. In contrast, using the bin divisions based on SAX algorithm, the entropy is 2.63 (Fig. 4A).

Leading motifs depending on interestingness measure. The 8 most interesting motifs obtained with use of every of the four measures of interestingness defined in Section 2.2.4, are presented in Fig. 5. An example of a feature vector based on these most interesting motifs is presented in Table 4.

A summary of all 20 (2 classes $\times$ 10 motifs) feature vectors obtained using the I_2 measure for all the data segments in our dataset (concatenated for the whole cohort) can be found in Supplementary Fig. 4. The hill-climbing method found $w = 15$, $\alpha = 10$, $k = 10$ and $I = I_2$ as parameters yielding the highest classification performance. Note that $k = 10$ was the highest allowed value in the hill-climbing setup. This suggests that segments of 625 [ms] in length and 10 bins per dimension should be used as an input to the SAX algorithm. Further, top 10 motifs should be delegated to represent each class and since I_2 was the leading method for quantifying interestingness, motifs should be chosen primarily due for their frequency, rather than due to their length or diversity. Also, performance accuracy is robust concerning α , w and k (Fig. 4B).

Comparison with other methods. The full summary of the performance of different classifiers can be found in Table 5. The new quantification method was able to achieve an average score of 0.86, and as such, it outperforms all of the other classifiers. Note that although the average classification score is the highest, not all the individual scores are the highest. In particular, the MultiLayer Perceptron classifier did not yield a high score for the k -motif method (0.69), while T-patterns scored high in that case (0.84). The highest score of any method-classifier combination was achieved by the k -motif method using k -Nearest Neighbors (0.94).

In this classification study, the validation was performed using stratified 10-fold cross validation (implemented with use of *sklearn* Python package). According to our results, the k -motif approach is significantly better than using full datasets, and using means and variances from raw experimental variables, but not significantly better than either the behavioural method²⁶ or the T-pattern method⁸. Significant differences between methods are visually presented in Fig. 4C (significance was assessed using one-sample t-test at confidence level of $p = 0.05$, with Bonferroni correction for multiple comparisons).

Discussion

The current approach examined the utility of the k -motif approach as a means to improve automated behavioural analysis in the open field. Our results suggest that this is an efficient technique which returns a sparse set of behavioural patterns best discriminating between experimental groups and gives higher classification performance than the available methods.

Figure 4. Results. (A) Proportion of visits to each bin after using SAX with parameters $\alpha = 5$ and $w = 15$. (B) Robustness of the k -motifs algorithm using k Nearest Neighbors using I_2 definition of interestingness, concerning parameters α , w and k . (C) Significance of differences between method performances. Yellow fields denote significant differences in performance between two corresponding methods, while purple means non-significant. Significance was assessed using one-sample t -test at confidence level of $p = 0.05$, with Bonferroni correction for multiple comparisons.

Normalising bins. The results were contrary to expectations, meaning that dividing the table into linear zones gives better performance than normalising divisions in creating a uniform distribution of zone visits. However, it is unclear whether this effect is specific to this particular dataset and additional validation with other datasets would be useful. While the literature suggests that normalising bins should aid in creating uniform visit distributions, it would be interesting to investigate the method further in other datasets.

Motifs, feature vectors and interestingness. The motifs yielded by k -motifs were interesting, although they may be difficult to interpret systematically. Most probably, motifs are specific to the particular shape and orientation of the open field arena used and the associated objects, and in another experimental setup, another set of motifs might come out as most predictive of the subject class. The k -motif approach, in the form proposed in this work, is an exploratory technique but is adaptable to different contexts. One noticeable property of motifs obtained in this study is that most of them, describe mostly linear motion or repetitions of linear motions. For example, many of the motifs in Fig. 5 show a single oscillation of a repeating motion, or else they show an almost linear movement. These movements are not very complex, and may, perhaps, be also picked up by a more constrained and computationally cheaper method. Some of the features of the motifs presented in Fig. 5 are particularly interesting. Of note, the motifs obtained from interestingness I_1 and I_2 are identical. Since the difference between these measures lies in the way that length and diversity of motifs are valued, while the value put on motif frequency is constant, this suggests that even in measure I_1 (the equal-weight measure) the frequency has more influence on the final selection. It may be useful to normalise the three properties before combining them into one measure, so that any of them has equal range of values, before taking the sum.

Furthermore, the nature of each measure of interestingness is reflected in the presented motifs. Both the measure focusing on length and the measure focusing on diversity do produce longer motifs, while the measure focusing on diversity additionally finds motifs with a larger number of unique values. Interestingly, multiple motifs are shared between different measures of interestingness focusing on length I_3 and diversity I_4 , while seemingly no motifs are shared between these and the measure focusing on frequency I_2 . This can be explained by the fact that frequency of motifs should be inversely proportional to both motif length and diversity. It should be noted that the most-interesting motifs presented in Fig. 5 include motifs of all relation types (positions and distances) except for the 'relative position' type. This means that no motif of this type was among the 8 most interesting based on any of the four interestingness measures. This may be explained by the fact that the relative position may have the highest variance of all of the spatial relations, meaning that patterns are less likely to occur.

In this study, I_2 , the measure focusing on frequency, increases classifier performance results more than measures based on length (I_3) and/or diversity (I_4). This might be associated with the fact that we are investigating quipirole model of the obsessive-compulsive behaviour. In the future research, it might be investigated whether the measures focusing on frequency (I_2) also to the highest classification performance in paradigms based on drugs with a different mechanism of action, e.g. repeated checking behaviour has also been associated with serotonin 5HT2C ligands²⁷.

Furthermore, in our study, repetitive and stereotypical behaviours indeed were revealed as the most informative features, which is expected given that we studies biomarkers of obsessive-compulsive behaviours. We cannot hypothesise about possible outcomes of this research when applied to other cohorts representing animal models of other cognitive disorders. Indeed, in paradigms in which we expect decreased general mobility and we do not anticipate any specific patterns of behaviour (such as models of depression), the outcome set of motifs would probably be more simplistic - but this is a subject to further research.

Figure 5. 8 Most interesting motifs resulting from the four different combinations of the three measures of interestingness as defined in Section 2.2.4. Note that one-dimensional features (distances) are presented in time scale, while two-dimensional features (movement) are shown in space. Relative movement patterns originate at the central position (0, 0). Where possible, bins (i.e. zones) are indicated with red lines (this is not possible for relative movement as the bins re-orient with each step).

29	22	5	19	5	21	2	21	17	16	22	29	19	21	21	17	13	18	17	20
----	----	---	----	---	----	---	----	----	----	----	----	----	----	----	----	----	----	----	----

Table 4. An example feature vector resulting from the quantification method, describing a single data segment. This vector contains 20 features, since there are two classes and $k = 10$, so 10 motifs from each class are used. Each number indicates the number of times a particular motif occurs in the data segment of interest.

Method	Mean	Gaussian Naive Bayes	Decision Tree	Multilayer Perceptron	k-Nearest Neighbors
Full Data	0.64 ± 0.16	0.72 ± 0.22	0.70 ± 0.14	0.56 ± 0.16	0.59 ± 0.13
Variable means and variances	0.68 ± 0.13	0.71 ± 0.15	0.66 ± 0.12	0.66 ± 0.16	0.68 ± 0.10
Behavioral method ²⁶	0.76 ± 0.12	0.77 ± 0.14	0.78 ± 0.10	0.74 ± 0.14	0.78 ± 0.09
T-patterns	0.83 ± 0.16	0.80 ± 0.13	0.88 ± 0.15	0.84 ± 0.16	0.81 ± 0.20
k-motifs	0.86 ± 0.10	0.88 ± 0.07	0.87 ± 0.09	0.69 ± 0.17	0.94 ± 0.05
Mean	0.76 ± 0.13	0.77 ± 0.15	0.78 ± 0.12	0.71 ± 0.14	0.76 ± 0.12

Table 5. Classification performances of the different quantification methods for each of the considered classifiers.

Supplementary Fig. 4 demonstrates some interesting properties of the feature vectors. The sessions are separated by class along the vertical axis, and a clear separation is visible in the values of the feature vectors. This visually shows the source of the classification accuracy as found in the evaluation. However, the main difference between the two groups seems to be that many motifs are frequent in the quinpirole treated group, while few motifs are frequent in the control group. Although this facilitates classification, it is contrary to expectations: since 10 motifs were taken from the top motifs in each class, the expectation would be for half of the motifs to be frequent in all sessions of one class, and the other half in the other class. One possible explanation for this phenomenon is the fact that motifs derived from the control class, although common in the control class, are even more frequent in the quinpirole treated class.

Parameters. Parameters optimized in the hill climbing algorithm specified twice as high number of bins per dimension, than it is specified in the T-pattern method. It would be interesting to explore how the T-pattern method would perform if using a higher number of bins. Although $k = 10$ was the optimum value, this was also the highest allowed value for k in the chosen parameter range. If higher values were allowed, the performance would likely keep increasing with k , since an increase in the number of features can only increase the classification performance. However, a higher value of k would also negatively impact the computational complexity of the algorithm. In addition to the four aspects of the method that were parameterised, some more possibilities exist that were not explored, e.g., a choice of which spatial relations to include. Also, interestingness of a motif could be defined in multiple other ways.

Evaluation. Our results indicate that - although the k -motif method performs significantly better than methods based on full data, and data compressed to means and variances - it does not perform significantly other than either of the two other, more complex methods. However, it should be taken into consideration that the parameters of the k -motif method were fine-tuned to our dataset, while the other methods were not privileged in a similar fashion. This manifests in the fact that the performance on the held-out data was only 0.75. The performance was lowest for the Multilayer Perceptron classifier, which is known to be sensitive to overfitting^{28,29}, which could also contribute to this effect. In the future, cross-validation of the results on another experimental cohort might be performed, although each cohort behaves differently in the open field environment and the classification performance would certainly drop.

In rodent studies in the open field, there are multiple confounding variables. Firstly, there are many different open field designs, and the items placed on the table can look and smell in various ways. Even if the objects are identical, the orientation of the table in space can matter to rodents, as the rodents are very sensitive to geographical directions³⁰. Locomotion patterns in rodents also depend on the day-night cycle³¹. Furthermore, rodents bred in different local conditions, fed with a different fodder, and in contact with different experimenters while maturing, can potentially exhibit different behaviours in an open field. These confounding factors cannot be fully mitigated.

For this reason, one can expect that some of the leading k -motifs will overlap between various open field experiments while some other motifs could differ from experiment to experiment. It would be very valuable to conduct a follow-up, comparative study in which the method is applied to a few experimental groups, and the results are compared.

One aspect of classification that was not explored in this project, but which would be very interesting for the future, is the evaluation of the relative contributions of features in the process of classification. It would be interesting to know which motifs, or which spatial relations, contribute most to determining classes in order to refine the k -motif algorithm further.

Conclusion

In this manuscript, we propose a novel concept for a data-driven quantification of behaviour with the use of k -motifs approach. We further apply this method to the dataset collected from an open-field experiment in the quinpirole rate model of obsessive-compulsive behaviour and demonstrate that the frequency-based motifs in behaviour lead to classification performance highly outperforming currently used methods for quantifying behaviour in the open field. In the future, the method could also be validated using cohorts of rodents exposed to a different experimental manipulation, e.g.:

1. NMDA antagonists such as phencyclidine and MK-801, which are known to produce complex behaviours in rodents such as stereotyped circling³²,
2. amphetamines (DAT blockade and VMAT transport reversal³³),
3. genetic rodent strains which are known to present with altered stereotyped behaviours such as self-grooming, e.g. BALB/c mice, BTBR mice, VDR knockout mice, Shank2 knockout mice, RLA rats.

Furthermore, our results suggest that taking data-driven approaches to the quantification of movement in the open field, is a good general direction for future research in this area. However, as there are multiple methods for finding patterns in the time-series data^{34,35}, and k-motifs is one of many possibilities, in the future, other possible approaches to detecting patterns in the behavioural recordings might also be explored.

Our method is also well-suited for assessing stereotypical and repetitive behaviours. One of such behaviours often studied in animal models of cognitive disorders, is patterns of locomotion from home base and back. Home base is defined as the place in which the animal stays for the longest cumulative time³⁶. For instance, home base is the pivotal point of many animal models of depression¹. In further research, we might consider a study in which we focus solely on motifs which contain the home base, and investigate whether in certain behavioural paradigms, these motifs are predictive for the disorder.

Data availability

All the codes developed for this manuscript, are publicly available at <https://github.com/MareinK/kmotifs-paper-code>. The behavioural datasets are handled by the TACTICS consortium and are not publicly available, and not published. For access to this data, please contact Dr. Jeffrey Glennon at jglennon77@outlook.com.

Received: 6 August 2019; Accepted: 12 December 2019;

Published online: 27 December 2019

References

1. Nestler, E. J. & Hyman, S. E. Animal models of neuropsychiatric disorders. *Nature Neuroscience* **13**, 1161 (2010).
2. Stewart, A. M. & Kalueff, A. V. Developing better and more valid animal models of brain disorders. *Behavioural Brain Research* **276**, 28–31 (2015).
3. Prut, L. & Belzung, C. The open field as a paradigm to measure the effects of drugs on anxiety-like behaviors: a review. *European Journal of Pharmacology* **463**, 3–33 (2003).
4. Szechtman, H. *et al.* Compulsive checking behavior of quinpirole-sensitized rats as an animal model of Obsessive-Compulsive Disorder (OCD): form and control. *BMC Neuroscience* **2**, 4 (2001).
5. Delcourt, J. *et al.* Comparing the EthoVision 2.3 system and a new computerized multitasking prototype system to measure the swimming behavior in fry fish. *Behavior Research Methods* **38**, 704–710 (2006).
6. Noldus, L. P., Spink, A. J. & Tegelenbosch, R. A. EthoVision: a versatile video tracking system for automation of behavioral experiments. *Behavior Research Methods* **33**, 398–414 (2001).
7. Heiderstadt, K., McLaughlin, R., Wrighe, D., Walker, S. & Gomez-Sanchez, C. The effect of chronic food and water restriction on open-field behaviour and serum corticosterone levels in rats. *Laboratory Animals* **34**, 20–28 (2000).
8. Casarubrea, M. *et al.* T-pattern analysis for the study of temporal structure of animal and human behavior: a comprehensive review. *Journal of Neuroscience Methods* **239**, 34–46 (2015).
9. Magnusson, M. M. Discovering hidden time patterns in behavior: T-patterns and their detection. *Behavior Research Methods, Instruments, & Computers* **32**, 93–110 (2000).
10. Salah, A. A., Pauwels, E., Tavenard, R. & Gevers, T. T-patterns revisited: mining for temporal patterns in sensor data. *Sensors* **10**, 7496–7513 (2010).
11. Szechtman, H., Sulis, W. & Eilam, D. Quinpirole induces compulsive checking behavior in rats: a potential animal model of obsessive-compulsive disorder (OCD). *Behavioral Neuroscience* **112**, 1475 (1998).
12. Lonardi, J. & Patel, P. Finding motifs in time series. In *Proceedings of the 2nd Workshop on Temporal Data Mining*, 53–68 (2002).
13. Li, Y. & Lin, J. Approximate variable-length time series motif discovery using grammar inference. In *Proceedings of the Tenth International Workshop on Multimedia Data Mining*, 10 (ACM, 2010).
14. Lin, J., Keogh, E., Lonardi, S. & Chiu, B. A symbolic representation of time series, with implications for streaming algorithms. In *Proceedings of the 8th ACM SIGMOD workshop on Research issues in data mining and knowledge discovery*, 2–11 (ACM, 2003).
15. Nevill-Manning, C. G. & Witten, I. H. Identifying hierarchical structure in sequences: A linear-time algorithm. *Journal of Artificial Intelligence Research* **7**, 67–82 (1997).
16. Lin, J., Keogh, E., Patel, P. & Lonardi, S. Finding motifs in time series. In *Proceedings of the 2nd Workshop on Temporal Data Mining*, at the 8th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (2002).
17. Lin, J., Keogh, E., Lonardi, S. & Chiu, B. A Symbolic Representation of Time Series, with Implications for Streaming Algorithms. In *Proceedings of the 8th ACM SIGMOD Workshop on Research Issues in Data Mining and Knowledge Discovery* (2003).
18. Russell, S. J. & Norvig, P. *Artificial Intelligence: A Modern Approach* (2nd ed.) (Upper Saddle River, 2003).
19. Pedregosa, F. *et al.* Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research* **12**, 2825–2830 (2011).
20. Murphy, K. P. Naive bayes classifiers. *University of British Columbia* (2006).
21. Safavian, S. R. & Landgrebe, D. A survey of decision tree classifier methodology. *IEEE transactions on Systems, Man, and Cybernetics* **21**, 660–674 (1991).
22. Gardner, M. W. & Dorling, S. Artificial neural networks (the multilayer perceptron): a review of applications in the atmospheric sciences. *Atmospheric Environment* **32**, 2627–2636 (1998).
23. Altman, N. S. An introduction to kernel and nearest-neighbor nonparametric regression. *The American Statistician* **46**, 175–185 (1992).
24. Van Asch, V. *Macro-and micro-averaged evaluation measures* (2013).
25. Diamantidis, N., Karlis, D. & Giakoumakis, E. A. Unsupervised stratification of cross-validation for accuracy estimation. *Artificial Intelligence* **116**, 1–16 (2000).
26. Zwamborn, J. *Feature validation in behavioural data*. Bachelor thesis, Radboud University, Nijmegen (2017).
27. Bailey, M. R. *et al.* The effects of pharmacological modulation of the serotonin 2C receptor on goal-directed behavior in mice. *Psychopharmacology (Berl)* **233**, 615–24 (2016).
28. Avrutskiy, V. Avoiding overfitting of multilayer perceptrons by training derivatives. *arXiv preprint arXiv:1802.10301* (2018).
29. Lawrence, S., Giles, C. L. & Tsoi, A. C. Lessons in neural network training: Overfitting may be harder than expected. In *AAAI/IAAI*, 540–545 (Citeseer, 1997).

30. Dudchenko, P. A. & Davidson, M. Rats use a sense of direction to alternate on T-mazes located in adjacent rooms. *Animal Cognition* **5**, 115–8 (2002).
31. Bains, R. S. *et al.* Assessing mouse behaviour throughout the light/dark cycle using automated in-cage analysis tools. *Journal of Neuroscience Methods* **300**, 37–7 (2018).
32. Ali, S. F., Newport, G. D. & Bracha, H. S. Phencyclidine and (+)-MK-801-induced circling preference: correlation with monoamine levels in striatum of the rat brain. *Neurotoxicol Teratol* **16**, 335–42 (1994).
33. Insel, T. R., Hamilton, J. A., Guttmacher, L. B. & Murphy, D. L. D-Amphetamine in obsessive-compulsive disorder. *Psychopharmacology* **80**, 231–5 (1983).
34. Mueen, A., Keogh, E. J., Zhu, Q., Cash, S. & Westover, M. B. Exact Discovery of Time Series Motifs. In *Proceedings of the SIAM International Conference on Data Mining* (2009).
35. Mueen, A. Time series motif discovery: dimensions and applications. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery* **4**, 152–9 (2009).
36. Eilam, D. & Golani, I. Home base behavior of rats (*rattus norvegicus*) exploring a novel environment. *Behavioural Brain Research* **3**, 199–211 (1989).

Acknowledgements

We would like to thank the members of the Translational Psychiatry group (PI: Jeffrey Glennon) for sharing their expertise on behavioural paradigms in rodents, and to the Department of Artificial Intelligence, Radboud University Nijmegen, for facilitating this research as a topic for MK. We would also like to thank Jesse Zwamborn, within the Radboud University AI department for his contribution to this project at the early stages. The research leading to these results was supported by the European Community's Seventh Framework Programme (FP7/2007–2013) under grant agreement no 305697 (OPTIMISTIC), the European Community's Seventh Framework Programme (FP7/2007–2013) under Grant Agreement No. 278948 (TACTICS) and European Union's Seventh Framework Programme for research, technological development and demonstration under grant agreement no 603016 (MATRICS).

Author contributions

M.K., N.B., K.K. and J.G. designed the study. K.K. performed the open field experiments and was consulting the team with respect to standard methods for quantifying behaviour in the open field. M.K. and N.B. conducted the computational study and drafted the manuscript. N.B., M.B. and J.G. co-supervised the M.K.'s Bachelor Thesis project. J.G., M.B., T.C. and J.B. critically revised the draft of the manuscript.

Competing interests

The authors declare that the research was conducted in the absence of any financial and non-financial relationships that could be construed as a potential conflict of interest.

Additional information

Supplementary information is available for this paper at <https://doi.org/10.1038/s41598-019-56482-z>.

Correspondence and requests for materials should be addressed to N.Z.B.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2019