Intelligent Human Systems Integration (IHSI 2025), Vol. 160, 2025, 678–686
https://doi.org/10.54941/ahfe1005874

Decoding Pet Signals: Bridging the Gap for Accurate Diagnoses

Rafael de Pinho André, Gustavo Tironi, Walléria Simões Correia, and Helouise Mattjie

Fundação Getulio Vargas, Rio de Janeiro, RJ 22253900, BRAZIL

ABSTRACT

The global pet care market is booming, with pet owners increasingly seeking ways to understand and improve the well-being of their animal companions. This demand extends beyond domestic pets to working animals, such as those in law enforcement and the military, and even to livestock in the agriculture industry. Accurate and timely diagnosis of health issues is crucial in all these sectors, but traditional methods often rely on subjective observations and infrequent veterinary visits. This paper introduces a novel Internet of Things (IoT) platform designed to bridge this gap by providing continuous, objective monitoring of animal behaviour and activity. Our platform employs a non-invasive, wearable device equipped with an array of sensors that capture physiological and movement data. The collected data is visualized through an intuitive dashboard accessible to pet owners, trainers, and veterinarians. This dashboard provides valuable insights into the animal's daily routines, activity levels, and potential anomalies. For pet owners, this translates to a deeper understanding of their pet's needs and early detection of potential health concerns. For trainers, the platform offers data-driven feedback to optimize training programs and monitor progress. Veterinarians can leverage the platform to access objective data, aiding in diagnosis and treatment planning, and enabling remote monitoring of patients. This research details the development and validation of the IoT platform, including the sensor technology, machine learning models, and dashboard design. We present results from a pilot study demonstrating the platform's effectiveness in accurately classifying animal activities and identifying deviations from normal behaviour patterns. The potential applications and implications of this technology are discussed, highlighting its contribution to improving animal welfare across various domains, from enhancing the bond between pets and owners to revolutionizing animal healthcare in veterinary practice and the agriculture industry.

Keywords: Behaviour classification, Wearable sensors, Smart monitoring, Internet of Things, Machine learning

INTRODUCTION

The market for canine activity monitoring devices has experienced rapid growth in recent years, fuelled by pet owners increasing awareness of and desire to improve their dogs' health and well-being (Grand View Research, 2024). These devices have gained popularity for their ability to track basic behaviours and physiological parameters, providing valuable insights into a dog's overall health.

However, there is a growing need for systems that offer more precise and detailed information about canine movement patterns. By leveraging advanced sensor technology and analytical techniques, these systems could enable the early detection of subtle behavioural changes, potentially aiding in the diagnosis of diseases and optimization of injury recovery protocols (Ladha et al., 2013). Emerging research has demonstrated the potential of such devices to address a range of health issues in dogs, including psychological disorders like separation anxiety (Wang et al., 2022) and chronic conditions such as epilepsy (Muñana et al., 2020 and Folkard et al., 2023).

Current canine activity monitors predominantly rely on 3-axis accelerometers to capture movement data (Chambers et al., 2021 and Uijl et al., 2017). Some devices incorporate additional sensors, such as gyroscopes (Ferdinandy et al., 2020 and Hussain et al., 2022) and, less commonly, magnetometers (Marcato et al., 2023). The combination of these three sensors forms an Inertial Measurement Unit (IMU), providing a reliable, lightweight, cost-effective, and energy-efficient solution for motion tracking. While alternative approaches utilize camera-based systems (Tsai et al., 2020 and Mealin et al., 2016) or even integrate sensor and camera data (Kim et al., 2022), this study focuses specifically on IMU-based systems for canine behavioural monitoring.

To contribute to this evolving field, we present a low-cost wireless wearable device constructed from commercially available hardware. Designed with an open-source framework, this device represents the first step in developing a novel Internet of Things (IoT) platform aimed at empowering veterinarians and pet owners to monitor daily routines, activity levels, and potential anomalies in dogs. This platform promises to provide valuable insights into canine behaviour, facilitating proactive health management and strengthening the bond between humans and their canine companions.

SETTING UP THE IoT WEARABLE DEVICE

On this Section, we summarize the assembling of the IoT device used to collect behavioural data during the experiment.

The IoT wearable device is a custom 3D-printed support and protective case that houses the microcontroller and its sensors. The custom 3D-printed support and protective case was designed to ensure secure and stable placement of the M5StickC Plus device. This support was then securely attached to a commercially available dog collar, allowing for comfortable and consistent positioning of the sensors around the dog's neck. This minimized sensor movement during the experiment, leading to the collection of reliable inertial data.

The M5StickC Plus device from M5 Stack, a commercially available platform for IoT development, was employed to collected inertial data. These devices contain a three-axis accelerometer with a sensitivity range of –16g to 16g and a three-axis gyroscope with a range of –2000 to 2000 degrees per second (DPS). This allowed for the capture of both linear acceleration and angular velocity, providing comprehensive information about the animal's movement and positioning.

The collected data was transmitted wirelessly to a remote database via Wi-Fi. This facilitated continuous data logging and allowed for real-time monitoring of the dog's activity. The use of a remote database also enabled efficient data storage and management, facilitating subsequent analysis.

To promote transparency and reproducibility, both the software developed for this study and the hardware schematics have been made publicly available as open-source resources. This allows other researchers and developers to utilize and build upon our work, fostering collaboration and accelerating advancements in the field of animal activity monitoring.

DATA COLLECTION

On this Section, we summarize the data collection procedure employed in the experiment. The correct positioning of the collar on the dog was verified before, during and after each session, to ensure consistent placement for reliable data capture and minimize variability due to sensor orientation.

Data Collection Procedure

The data collection process was structured into a series of distinct experimental trials to facilitate data synchronization and labelling. Each trial adhered to the following workflow:

- Initiate video recording of the dog's activity;
- Perform the synchronization procedure;
- Capture specific behaviour;
- Conclude video recording of the dog's activity.

To ensure manageable video lengths and optimize labelling accuracy, each recording was kept under 120 seconds. Each trial focused on eliciting a specific set of behaviours to maintain control and facilitate the annotation process.

The synchronization procedure aimed to establish a temporal alignment between the video recording and the inertial sensor data. With the video recording initiated and the device held steady, the researcher focused their attention on a precision clock to capture the displayed time. This temporal reference point, clearly visible in the video frame, was crucial for subsequent synchronization of the video and inertial sensor data. Following the capture of the time, the researcher performed two distinct movements:

- Sharp Agitation: The device was abruptly shaken for approximately 1–2 seconds;
- Sudden Stop: The agitation was stopped abruptly, and the device was held motionless for 1–2 seconds.

After completing these movements, the recording continued without interruption, capturing the planned behaviours for the trial.

Behavioural Categories and Elicitation

The following predefined behavioural categories were used for annotation, along with the corresponding elicitation techniques:

- Walking: No specific trial was dedicated to walking. This behaviour was captured during other trials as the dog moved naturally between activities. Walking was characterized by a slow, coordinated movement of the limbs, with the legs moving in a specific sequence (left hind leg, left front leg, right hind leg, right front leg), resulting in forward motion.
- Sitting: The dog was cued to sit between 2 and 4 times during the trial, with each sitting instance lasting approximately 10 to 15 seconds. Sitting was defined as the dog having all four extremities and rump in contact with the ground, exhibiting minimal to no movement.
- Jumping: The dog was encouraged to perform between 4 and 5 jumps during the trial. Jumping was characterized by the dog propelling itself upward, with all four paws leaving the ground, and landing on a higher surface. This behaviour was often elicited in conjunction with "Descending from an Object."
- Descending from an Object: The dog was guided to descend from an object, such as a sofa, between 4 and 5 times during the trial. This behaviour was typically paired with the "Jumping" task.
- Laying Down: The dog was cued to lay down between 3 and 4 times during the trial, maintaining the position for 15 to 20 seconds each time. Laying down was defined as the dog's torso being in contact with the ground, with hips and shoulders at the same level. Variations, such as laying on the side or back with legs in the air, were also included in this category.
- Running: The dog was encouraged to run for approximately 10 seconds, between 2 and 3 times during the trial. Running was characterized by a coordinated and rapid movement of all four limbs, with all four paws simultaneously leaving the ground at some point during the stride. Galloping, a specific type of running gait, was typically observed during playful activities.
- Standing: No dedicated trial was conducted for standing. This behaviour was recorded during other trials when the dog was stationary. Standing was defined as the dog having all four extremities on the ground, without the torso touching the ground, and exhibiting minimal to no movement.

Experimental Trials

Based on these behavioural categories, three distinct experimental trials were conducted, each recorded in a separate video and following the workflow:

- Sitting Trial: Focused on eliciting and capturing instances of sitting behaviour
- Jumping and Descending Trial: Focused on capturing jumping and descending behaviours.
- Laying Down Trial: Focused on eliciting and capturing instances of laying down behaviour.

The behaviours “Walking,” “Running,” and “Standing” were captured opportunistically during these trials. When necessary, additional trials were conducted to specifically elicit these behaviours.

Behaviour Annotation

During offline data analysis, video recordings were meticulously synchronized with the inertial data streams, adhering to the established protocol. This synchronization ensured accurate alignment between the observed behaviours and the corresponding sensor readings. Each video frame was then carefully annotated with predefined labels representing distinct motion patterns in the dog’s behaviour, such as walking, trotting, running, playing, and resting. This frame-by-frame annotation process was facilitated by Label Studio, a commercially available software commonly used for subtitle editing, which proved to be an effective tool for this task.

To maintain objectivity and ensure the integrity of the dataset, only video segments where the dog’s activity could be clearly and unambiguously identified were labelled. This resulted in a dataset with interrupted labels, as periods of ambiguous or transitional movements were excluded. Additionally, any observed body postures that did not align with the predefined activity categories were also excluded from the analysis.

Following the video annotation, the raw IMU data, initially sampled at a higher frequency, was down sampled to 30Hz to match the frame rate of the video recordings. This down sampling process involved grouping consecutive sensor readings and calculating the average value for each group, effectively reducing the data resolution while preserving the overall motion trends. Once the IMU data was synchronized with the video frame rate, it was merged with the labelled dataset. This resulted in a comprehensive and cohesive annotated dataset containing both the raw inertial measurements and their corresponding activity labels, ready for subsequent analysis and model training.

DATA ANALYSIS

All classification models were developed using Python (v3.12.3), with the scikit-learn library (v1.5.2) being used to implement machine learning techniques. Two algorithms were selected for analysis: Random Forest (RF), k-Nearest Neighbours (KNN).

A grid search algorithm was employed to explore all possible combinations of the Random Forest classifier’s two hyperparameters: maximum depth (with values of 5, 10, and 15) and the number of estimators (set to 100, 250, and 500). For the KNN model, various weights and metrics were tested, with the value of k fixed according to the number of distinct behaviours present in the dataset, utilizing 5-fold cross-validation for model evaluation.

The evaluation was performed using an intra-subject approach, where training and testing were conducted on the same individual. This was done because the primary focus was not on evaluating the model’s generalization. We report average accuracy and F1-score, as these metrics provide a comprehensive evaluation of the model’s ability to correctly predict the class in multi-class classification problems with a natural class imbalance. To assess the model performance, we employed three different sets of features: accelerometer data alone, gyroscope data alone, and a combination of both. The results are shown in Table 1 and 2.

Table 1. Average classification results for Random Forest.

Random Forest	Accelerometer	Gyroscope	Accelerometer and Gyroscope
Accuracy	66.8	61	85.3
F1-Score	56	48	82

Table 2. Average classification results for KNN.

KNN	Accelerometer	Gyroscope	Accelerometer and Gyroscope
Accuracy	65.4	59.5	62.5
F1-Score	57	52	55

Figure 1: Confusion matrix.

Figures 1 and 2 show the confusion matrix alongside the distribution graph of behaviours, to allow for the evaluation of the error for each behaviour. The model exhibited a clear bias towards predicting "standing," significantly exceeding its true prevalence in the dataset. This over-representation is likely attributable to the inherent imbalance in the data, where "standing" constitutes the most frequent class. Preliminary attempts to address this imbalance using basic oversampling techniques proved unsuccessful, leading to a marked decline in the model's overall performance. This underscores the need for more sophisticated strategies to either balance the dataset or to implement a more controlled and balanced data collection protocol. However, achieving balanced data collection in this context presents unique challenges due to the non-collaborative nature of the canine subjects. This issue will be a primary focus of future research endeavours.

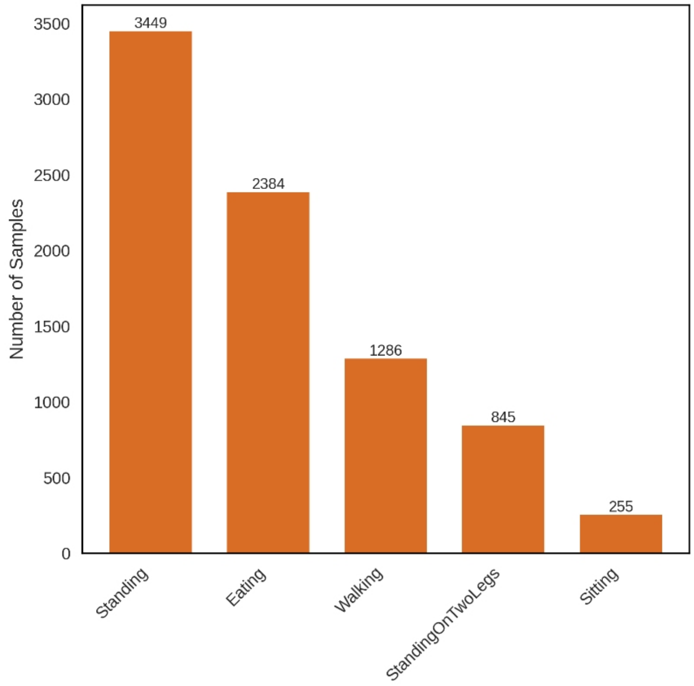

Figure 2: Distribution graph of behaviours.

The model frequently exhibits inconsistencies in its predictions within a single behavioural sequence. For instance, it might accurately classify ten consecutive frames as "standing" before misclassifying a single frame as "eating," only to revert to correctly classifying the subsequent twenty frames as "standing." Such abrupt transitions are highly implausible, as canine behaviour typically unfolds gradually and continuously. While behaviours naturally have a start and end point, these transitions occur over a period, not in discrete, instantaneous shifts.

This limitation stems from the model's treatment of individual frames as isolated instances, neglecting the inherent temporal continuity of behaviour. By disregarding the sequential nature of actions and their varying durations, the model generates predictions that lack realism and fail to capture the nuances of natural movement. While treating the data as windows, rather than individual frames, could potentially improve the model's performance, this approach necessitates careful consideration of the train-test split to avoid data leakage. Improper splitting could inadvertently introduce information from the test set into the training process, leading to artificially inflated performance metrics.

To mitigate this issue, several strategies could be explored. One potential solution involves incorporating the model's previous prediction as a feature, allowing it to leverage temporal context and better capture transitions between behaviours. Alternatively, a more fundamental shift in approach could involve designing a model that classifies behavioural sequences or regions, rather than individual instances. Such a model would explicitly define the start and end points of each behaviour, providing a more realistic and temporally consistent representation of how a dog's actions evolve over time. This would ultimately lead to more accurate predictions that reflect the natural flow and progression of canine behaviour.

CONCLUSION

This study aimed to validate the feasibility of collecting and utilizing inertial data from dogs using a low-cost, open-source, and easily deployable device. While the activity classification model achieved satisfactory performance, the primary focus was not to develop a state-of-the-art classifier but to demonstrate the potential of this approach for gathering rich and informative data on canine behaviour. The successful implementation of this system paves the way for future research and development in the field of canine activity monitoring.

Future work will focus on several key areas to enhance the platform's capabilities and generalizability. This includes expanding the dataset with more dogs and a wider range of breeds to improve the model's robustness and address potential breed-specific variations in movement patterns. Furthermore, exploring semi-supervised learning techniques could mitigate the challenges of manual labelling and potentially lead to the development of a standardized dataset for canine activity recognition. In addition to refining the classification model, future efforts will involve incorporating additional sensors, such as microphones, to capture a broader spectrum of behavioural cues. Finally, the development of an intuitive and informative dashboard will be crucial for translating the collected data into actionable insights for veterinarians and pet owners alike.

REFERENCES

Ali Hussain, Sikandar Ali, Abdullah, and Hee-Cheol Kim. Activity detection for the wellbeing of dogs using wearable sensors based on deep learning. IEEE Access, 10:53153–53163, 2022.

Bence Ferdinandy, Linda Gerencsér, Luca Corrieri, Paula Perez, Dóra Újváry, Gábor Csizmadia, and Ádám Miklósi. Challenges of machine learning model validation using correlated behaviour data: Evaluation of cross-validation strategies and accuracy measures. PLOS ONE, 15(7):1–14, 07 2020.

Cassim Ladha, Nils Hammerla, Emma Hughes, Patrick Olivier, and Thomas Ploetz. Dog's life: Wearable activity recognition for dogs. In Proceedings of the 2013 ACM.

Emily Folkard, Lee Niel, Luis Gaitero, and Fiona May Kier James. Tools and techniques for classifying behaviours in canine epilepsy. Frontiers in Veterinary Science, 10, 2023.

Grand View Research. Pet wearable market size, share & trends analysis report by technology (rfid, gps, sensors), by product, by animal type, by component, by application, by sales channel, by region, and segment forecasts, 2024–2030, 2024. Market Research Report.

Huasang Wang, Othmane Atif, Jirong Tian, Jonguk Lee, Daihee Park, and Yongwha Chung. Multi-level hierarchical complex behavior monitoring system for dog psychological separation anxiety symptoms. Sensors, 22(4), 2022.

Ingrid den Uijl, Constanza B. Gómez Álvarez, David Bartram, Yoni Dror, Robert Holland, and Alasdair Cook. External validation of a collar-mounted triaxial accelerometer for second-by-second monitoring of eight behavioural states in dogs. PLOS ONE, 12(11):1–13, 11, 2017.

International Joint Conference on Pervasive and Ubiquitous Computing, UbiComp '13, pages 415–418, New York, NY, USA, 2013. Association for Computing Machinery.

Jinah Kim and Nammee Moon. Dog behavior recognition based on multimodal data from a camera and wearable device. Applied Sciences, 12(6), 2022.

Karen R. Muñana, Julie A. Nettifee, Emily H. Griffith, Peter J. Early, and Nathanael C. Yoder. Evaluation of a collar-mounted accelerometer for detecting seizure activity in dogs. Journal of Veterinary Internal Medicine, 34(3): 1239–1247, 04 2020.

Marinara Marcato, Salvatore Tedesco, Conor O'Mahony, Brendan O'Flynn, and Paul Galvin. Machine learning based canine posture estimation using inertial data. PLOS ONE, 18(6):1–28, 06 2023.

Ming-Fong Tsai and Jhao-Yang Huang. Predicting canine posture with smart camera networks powered by the artificial intelligence of things. IEEE Access, 8:220848–220857, 2020.

Robert D. Chambers, Nathanael C. Yoder, Aletha B. Carson, Christian Junge, David E. Allen, Laura M. Prescott, Sophie Bradley, Garrett Wymore, Kevin Lloyd, and Scott Lyle. Deep learning classification of canine behavior using a single collar-mounted accelerometer: Real-world validation. Animals, 11(6), 2021.

Sean Mealin, Ignacio X. Domínguez, and David L. Roberts. Semi-supervised classification of static canine postures using the microsoft kinect. In Proceedings of the Third International Conference on Animal-Computer Interaction, ACI '16, New York, NY, USA, 2016. Association for Computing Machinery.