PetTrack: Tracking Pet Location and Activity Indoors

ABSTRACT

Tracking a pet's location and activity indoors is of peak interest to pet owners who want to feel connected to their pet and to pet owners who are concerned for their pet's well-being. Today, such tracking is performed using pet cameras. However, cameras need to be installed in every room accessible to the pet, do not work well in the dark, and generate a tremendous amount of data. In this paper we develop an indoor localization and pose detection system for pets using ultra-wideband (UWB) radios and an accelerometer affixed to the pet's collar. We also develop a new message format for UWB packets to carry accelerometer data. We create a pose-estimation logic to detect the pet's pose using the accelerometer data, and develop a functional mobile application that shows where the pet is and their activity. We believe this cost-effective and efficient new way of tracking a pet indoors will inspire others to extend this research further.

1 INTRODUCTION

Monitoring pets in indoor spaces is predominantly performed through pet cameras today. Indeed cameras offer direct view of our pets and pet-owners enjoy the connection they feel with their pets even when the owners are away from home. However, cameras are required to be placed in every room (and sometimes several cameras are required in each room), cameras consume a very high data bandwidth for streaming and storage, and cameras have known to cause privacy risks due to the likelihood of them being hacked. In this work, we ask these questions: *Is it possible to track a pet's activities using body-worn sensors and wireless localization instead of cameras?* In answering this question in the affirmative, we have created PetTrack, an ultra-wideband (UWB) localization solution for tracking a pet's indoor location, and monitoring the pet's activity. The system comprises of a set of UWB anchors installed in the indoor space, and a UWB device worn by the pet, on the collar or a harness. The UWB device is also equipped with an accelerometer to detect the pet's posture and infer activities. Our custom platform also allows adding more sensors to the collar, including microphones, whistles, etc. for better monitoring and management of the pet's activities.

But before we describe the platform and our system design for PetTrack, it is important to motivate the features enabled by PetTrack and its advantages over a camera-based solution. (1) PetTrack directly provides the coordinates of the pet inside an indoor space. This enables simple queries such as "how much time does my pet spend in the living room?" The analysis required to answer such a question is minimal compared to the analysis needed by a camera-based system. Thus, PetTrack provides a simpler, less compute-intensive solution to pet-location analytics. (2) PetTrack uses the orientation data measured from the collar device to deduce the pet's pose, which is crucial in exercise and weight management for pet physical therapy, and for monitoring general pet health. While we have not achieved fine-grained pose estimation, we are able to distinguish between sitting, standing, walking, and sleeping poses. Pet-pose analytics is thus greatly simplified compared to a camera-based system. (3) The data bandwidth required to transmit the pet's location and pose information over the network is quite small. This saves home Internet data usage and mobile cellular data usage when accessing the information. A vast majority of video data captured by cameras is never used. By distilling out only the most relevant information through location and orientation sensors, PetTrack takes a minimalist data approach. (4) Wireless signals can penetrate through walls, furniture, etc. Therefore a small number of UWB anchors suffice to cover the entire home. In contrast, camera coverage is limited to a single room, requiring several cameras to be installed in a typical home. Thus, PetTrack requires minimal hardware to track a pet.

In this work, the core principle for pet localization is wireless distance measurements. The ultra-wideband radio on the pet's collar performs wireless ranging with fixed anchors in the environment. The pet's location is solved using trilateration, which involves converting the time of flight between each anchor and the collar device to distances and using those distances to each anchor to determine the indoor location of the pet. The calculated location of the pet is then periodically sent to a cloud server, from where it is available to the pet's owners. At the same time, an accelerometer on the collar device piggy-backs its data on the UWB packets. Estimations of the pet's pose are made using the accelerometer readings, coupled with the pet's movement data from the wireless ranging. An aggregating compute device in the home, a raspberry pi in our prototype, performs the trilateration and the pose estimation. This aggregating device is connected to the Internet over the home network and keeps a cloud service informed about the pet's whereabouts.

Submitted to BodySys '22, June 25–July 01, 2022, Portland, OR
2022. ACM ISBN 978-x-xxxx-xxxx-x/YY/MM...$15.00
https://doi.org/10.1145/nnnnnnn.nnnnnnn

Our contributions are:

(1) We introduce a methodology for localization of a pet in an indoor environment using UWB ranging.
(2) We introduce a system that combines accelerometer data with UWB ranges to determine a pet's pose
(3) We introduce a mobile application that renders the pet's whereabouts and pose in real-time.

2 BACKGROUND ON SENSORS

PetTrack primarily relies on two sensors: (a) ultra-wideband radios for distance measurements, (b) inertial sensors for orientation measurements. In this section, we will describe background material on sensors themselves, including the standard ways of using them. In the next section, we will delve into our specific improvements to the state-of-the-art localization techniques.

Ultra-wideband Localization: The core idea in wireless localization is measurement of the amount of time it takes for wireless signals to go from one device to another. Since wireless signals travel at the speed of light (3×10^8 m/s), this time of flight for wireless signals must be computed at nanosecond precision. A challenge in precise time of flight measurement is in obtaining exact arrival times. For a narrow-band signal, the arriving signal rises above the noise floor rather slowly, making it difficult to tell what the exact arrival time is. Hence, we need to use signals with large bandwidths—ultra-wideband signals. Such signals rise quickly above the noise floor within a few picoseconds and therefore result in highly precise arrival time estimates.

However, accurate estimation of arrival times, while necessary for precise distance measurements, is not sufficient. We also need tight synchronization between the clocks of the two devices involved in the message exchange. Intelligent ranging protocols have been devised to overcome this challenge. For example, the alternative two-way ranging [11] relaxes the constraints of response time through a crafted formulation that naturally removes the clock drift errors, allowing more flexibility to the participating devices and finding wide adoptions [7].

Inertial Sensors: Ego-centric tracking of the orientation can be performed using inertial sensors. Typically this includes accelerometers, gyroscopes, and magnetometers. Unfortunately, indoor spaces are not amenable to the use of magnetometers since various electrical and electronic components in a household frequently corrupt the magnetometer readings. Gyroscopic movements only measure the angular velocity, which in this context does not provide much relevant data without frequent re-calibration. Hence, we only focus on the use of accelerometer output.

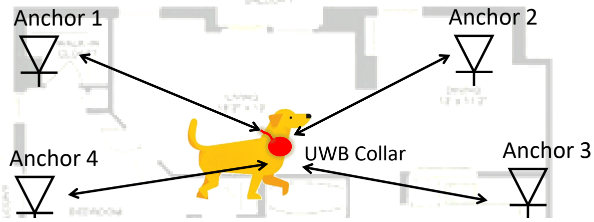

Figure 1: PetTrack System design includes several fixed UWB anchors in a home, with the pet wearing a UWB collar.

3 PETTRACK SYSTEM DESIGN

Ultra-wideband (UWB) localization techniques have existed for several years now. Therefore, it might seem that localization for any application can simply use the existing UWB ranging and trilateration methods. However, tracking a pet's location and capturing their pose will need several modifications to the standard two-way ranging scheme [5]. As shown in Fig. 1, we install a small number of anchor devices in the pet's home, and affix a client device to the pet's harness or collar. We now present the modifications made in PetTrack below.

3.1 One to Many TWR Ranging

Serially performing two-way ranging (TWR) with every anchor and the pet's UWB device is a slow process since every individual ranging operation takes time. We find that pipelined two-way ranging mitigates this issue by reusing just a single message transmitted by the UWB to behave as a trigger for *all anchors* to respond.

More specifically, in PetTrack, the pet's wearable device performs a pipelined two-way ranging[4] with the deployed anchors in the indoor environment. The pet's wearable device initiates TWR by sending a POLL message, which contains a schedule for the anchors' transmissions. Upon receiving this POLL message, the anchors transmit the RESPONSE messages in their respective slots dictated by the schedule. The pet's wearable device replies with a single FINAL message (see Fig. 3). The POLL, RESPONSE, and FINAL messages contain the receive and transmit timestamps necessary for calculating the Time-of-Flight(ToF) between the pet's wearable device and each anchor. Note that the schedule can be configured by the user, or through a one-time neighbor discovery protocol at system initiation. We program it based on the hard-coded anchor IDs.

3.2 Eavesdropper based Location Solver

Conventionally, in TWR, each anchor measures its ToF from the pet's wearable device, and all the ToF measurements are then collected by a central processor where a localization solver produces a location result. However, doing so requires

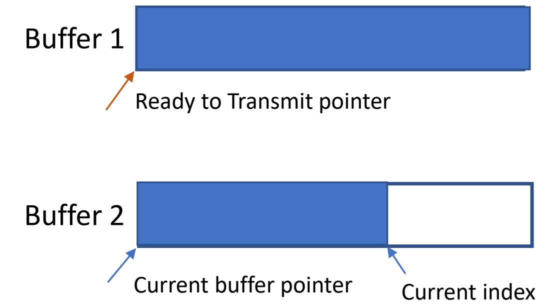

Figure 2: Illustration of Double Buffering

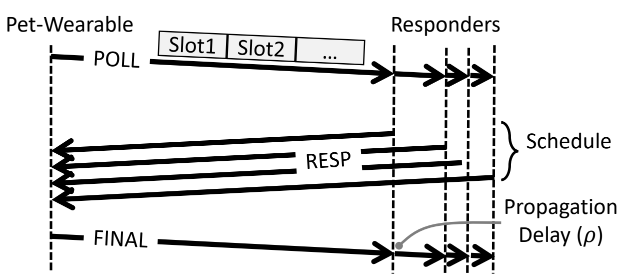

Figure 3: Pipelined TWR between the Pet's Wearable Device and Anchors.

additional hardware such as a wired backbone network, or dedicated ToF collection time slots, which reduce the localization rate. We propose an alternative approach where an eavesdropping device at the central processor overhears the TWR messages and computes the ToFs between the pet wearable device and the anchors. The only requirement is that all the receive and transmit timestamps at the anchors and the pet wearable device are embedded in the TWR messages, a requirement of the conventional TWR. However, in addition, the anchor's receive timestamp of the FINAL message also needs to be transmitted, which can be embedded in its next RESPONSE message. Once the ToF measurements are available at the eavesdropper, the central processing unit then computes the pet's location using a least-sqaure solver [10].

3.3 Accelerometer Data in UWB packets

In addition to performing TWR, the pet wearable device continuously collects the 3-axis acceleration data from the accelerometer on the pet-wearable device. To keep the computation complexity low, raw accelerometer data is transmitted to the central processing unit for pose classification. We exploit the existing TWR messages to carry the accelerometer data. To deal with the asynchronous nature of data collection and communication, we use a double buffering technique (see Fig. 2) to handle the sampled accelerometer data on the pet's wearable device: the data samples are stored in one of two buffers, and when it is filled, this buffer is marked as "Ready-to-Transmit" while the newer samples are stored in the other buffer; before transmitting the TWR messages, the pet wearable device checks the buffer status and will embed the entire data buffer in the TWR message if one of the buffer is marked as "Ready-to-Transmit". Note that both POLL and FINAL messages transmitted by the pet wearable device can carry the accelerometer data. This ensures the same length of accelerometer data is received by the pose classifier, making the processing simpler. The TWR message structure is shown in Fig. 4.

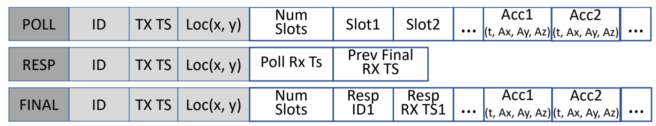

Figure 4: Packet structure for TWR message exchanges.

3.4 Pose Inferencing

Understanding the pose of the pet can help us gain insights into the activities that a pet performs throughout the day. To detect the pet's motion, an accelerometer has been integrated on the pet's wearable device. The 3-axis acceleration data is collected at $50Hz$ and is processed by the central device for pose classification. To visualize how the acceleration data can be used for pose inference, we show a 3D scatter plot(Fig. 5) from our measured data, where different poses manifest as distinct clusters.

Machine learning techniques have been widely used for classification tasks. For its simplicity and wide application, we use the K-nearest neighbor(KNN) classifier to let the system automatically determine the pose of the pet. The accelerometer data is taken in windows of a fixed size. Because of our double buffering technique, the accelerometer data is always taken at a constant rate and therefore doesn't contain any jumps in time. To remove the data related to movements and transitions between poses, we first check each windowed signal based on its variance to only retain the data that can be classified as **static poses** (with low variance). Then the windowed signal is vectorized into $[Ax_1, Ay_1, Az_1, ..., Ax_n, Ay_n, Az_n]$ and fed into the KNN classifier for classification.

3.5 Cloud Storage

To render read-time data on the app and persist it for later use for offline analysis we use Firebase as a cloud storage platform. Cloud storage is a common Cloud Computing model to store data on the Internet through various cloud computing service providers who manage and operate data storage as a service. For our work, we use Firebase which is a simple, powerful and cost-effective service built by Google to scale. To capture data in real time, we used Firebase Real-Time Database which stores data as a JSON and synchronizes it in

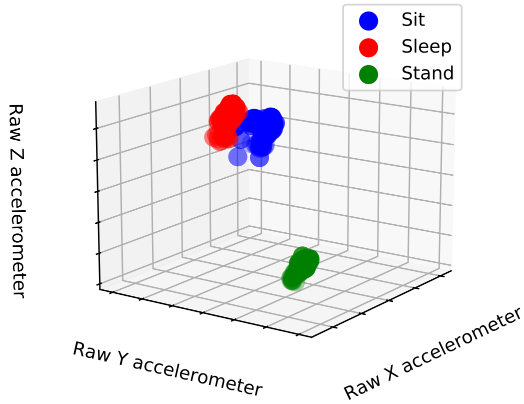

Figure 5: Visualizing the clusters of accelerometer measurements in different poses

real-time to every connected client. The most important reason why we use Firebase is that we can build cross-platform apps with iOS, Android, JavaScript and all of our clients share only one real-time database instance and automatically receive updates with the newest data.

3.6 Mobile Application

Since the focus of this work is on the techniques for localization of the pet, we have created only a minimal mobile application. Currently, the underlying floorplan map is fixed, and the anchor locations are manually inputted.

4 IMPLEMENTATION

We have created a custom-built PCB that houses a UWB chip, and an accelerometer. A cortex M0 microcontroller on the PCB runs the ranging code by interfacing with the Decawave DWM1000 UWB module and samples the accelerometer data from an ADXL335 chip using the microcontroller's ADC. This device is also provisioned with a piezoelectric buzzer, an SDCard, and a real-time clock. Code is uploaded to the device via a micro USB port. The device can be powered either through a micro-USB port or via a battery. We use copies of this device as the pet's wearable device where all of its sensors are utilized, as well as anchors mounted on the wall at various places in the house. Anchors only use the microcontroller controlled UWB module, and are powered via a USB charger. The pet's wearable device is powered using a 1200mAh LiPo battery. The eavesdropper is another copy of the same device, but only the UWB module is used as a gateway between the UWB data exchanges and the Raspberry Pi which finally computes the pet's location and pose information. The eavesdropper UWB device is directly connected with a RaspPi 4 and prints out all observed packets on its serial interface. The Rasp Pi captures this information via a USB-to-Serial driver and runs a python program to

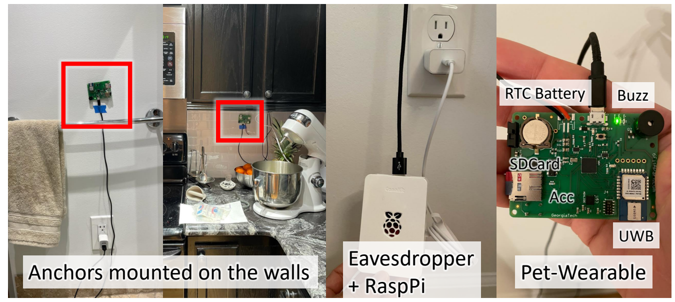

Figure 6: Our experimental setup using our custom wireless ranging and accelerometer devices.

continuously deduce the pet's location. The UWB packets originating from the pet's wearable device also include accelerometer readings which are processed at the Rasp Pi in real-time to infer the pet's pose. Fig. 6 shows how the anchors were mounted on walls of a two bedroom apartment, how the pet-wearable device was mounted on a pet dog(Fig. 7)[1], and the eavesdropper+Raspberry Pi combination which captures all exchanged data. Separating our the RaspPi from the pet-wearable device allowed us to keep the weight of the pet-wearable device to a minimum at about 17.886g.

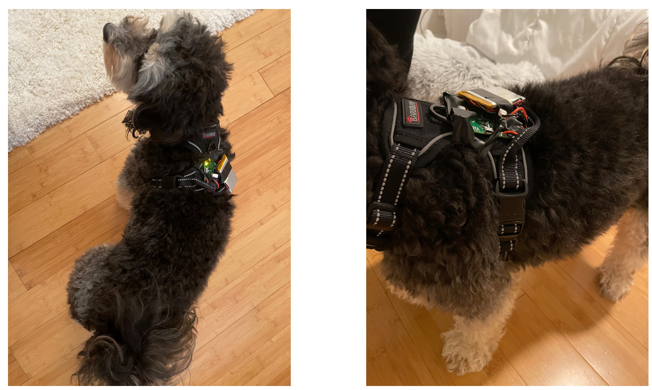

Figure 7: Pictures of the Pet's Wearable Device

Obtained information at the Raspberry Pi is transmitted to a Firebase cloud server. A python application on the eavesdropper's Raspberry Pi device collects all the localization and accelerometer data in one place. This data is obtained over the serial port by the Raspberry Pi. The placement of the anchors in the house is assumed to be known. To store the obtained data on cloud, we use Firebase. It's data config contains apikey, authodomain, databaseURL, and storageBucket. We push all the obtained raw information to the firebase database. An Android app obtains this information and converts it to an animated pet character superimposed

[1]This project has obtained the required permissions from both the institutional review board and the Institutional Animal Care and Use Committee

	Sit	Sleep	Stand
Sit	100%	0	0
Sleep	0	100%	0
Stand	0	6.9%	93.1%

Table 1: Normalized Confusion matrix of KNN pose classification

on the home's floor plan. The user is expected to provide the floorplan beforehand, and is also expected to mark the locations of the anchor devices on the floorplan. At this time, our mobile app does not allow marking the anchor locations on the mobile phone, but this feature is planned for the future.

Our implemented system runs at about $1 - 2\,Hz$ end-to-end given the delays introduced by the Firebase cloud server. Raw location estimates are produced at the RaspPi at a rate of $10\,Hz$ and we expect that a dedicated cloud service in the future will improve the end-to-end update rate.

5 EVALUATION

We evaluate PetTrack on a dog[2] of breed tiny bernedoodle in a 1350 sq.ft. apartment. The pet's wearable device was attached to the back of the harness behind the subject dog's neck. We evaluate the localization and pose inference accuracy with the subject dog being instructed to stay in static poses. Then we show a visualization of activity tracking of the subject dog freely moving around. The anchor locations are obtained using a floorplan and a laser ranger.

5.1 UWB Localization

The subject dog wears the harness equipped with the pet's wearable device, and is instructed to remain in a certain static pose at a certain location while localization data is recorded. In Fig. 8, the scatter plot shows a visualization of the localization results. For quantitative analysis, we compute the localization error and plot the cumulative distribution function (CDF) in Fig. 9 compared to the centroid tag location. The subject dog is in Sleep pose at P1-P2, Sit pose at P3-P4, and in Stand pose at P5-P7. Overall, at most locations while under difference poses, the 75^{th} percentile localization error is less than one meter.

5.2 Pose Inference

The subject dog is instructed to remain in different static poses (Sit, Sleep, Stand) during training and testing. The real-time accelerometer measurements are used to generate the predicted pose labels. The classification accuracy is 98.6%, and the confusion matrix is shown in Table 1. This result shows our classifier is reliable when the pet remains in the same pose.

[2]Our study has been approved by our institution's IACUC and IRB.

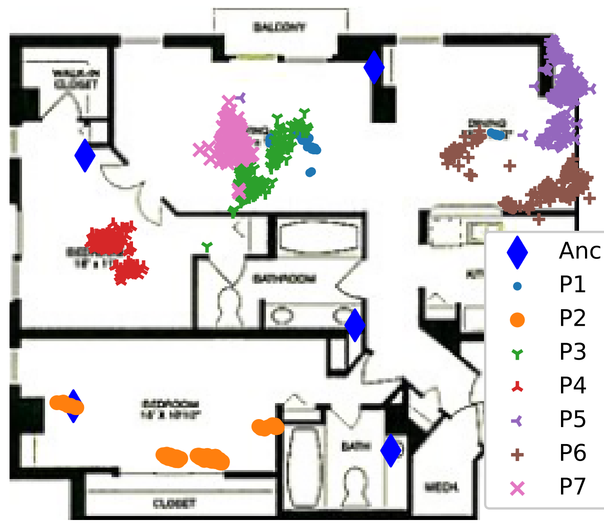

Figure 8: Scatter plot of localization results at 7 different pet locations (P1-P7)

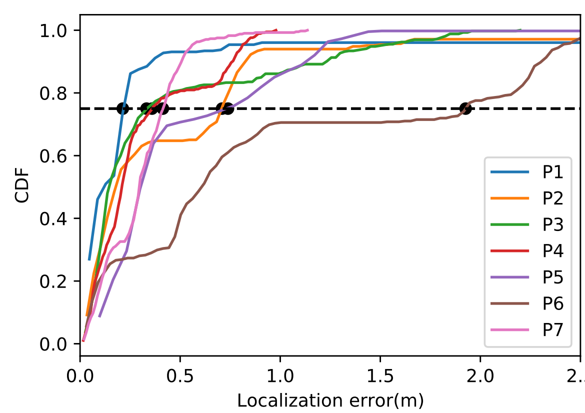

Figure 9: CDF of localization error at pet locations P1-P7.

5.3 Free Moving Experiment

In this experiment, the subject dog moves freely in the test environment, while the localization and pose inference data are recorded. The subject dog's movement is videotaped for obtaining its ground truth position and activity. To show the performance of the classifier, we do not consider the data where the ground truth is non-static. Fig. 10 shows the inferred pose and the ground truth pose over time. The classification accuracy is 88.4%.

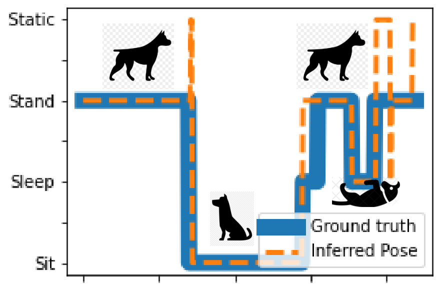

Figure 10: Free Moving pose inference over time

6 RELATED WORK

The area of pet analytics is an upcoming field. This is also reflected by the fact that pet ownership is on the rise—71% of households have pets in 2021 compared to 67% in 2019[1].

Outdoor localization of pets has been enabled using a GPS tracker on the pet's collar for tracking lost pets [13]. However, GPS based solutions are not reliable for indoor environment. In recent years, indoor localization solutions [6, 15] have shown tremendous success, especially the UWB systems [4, 12], achieving decimeter-level accuracy. UWB tracking has seen applications in first responders [4], robots [8], and mobile phones [9]. But using indoor localization for tracking and understanding pets' activities have been missing in literature.

Besides location tacking, pose tracking is also crucial for understanding the pets' activity and overall health. Camera and convolutional neural network (CNN) based pose tracking has enabled dog pose recognition and reconstruction [14]. However, continuously running a CNN on live visual data is very computationally expensive and takes a lot of processing power, and the accuracy and reliability of the pose estimation is highly dependent on the placement of cameras and object occlusions. Accelerometer based pose recognition is a much more cost effective, data efficient method that have demonstrated high classification accuracy [2, 3], which inspired us to adopt accelerometer as well.

7 DISCUSSION AND FUTURE WORK

How different is pet tracking from tracking people? Typically, when people use localization technology, they wish to track their own location. However, in pet tracking, pet-owners need to know the location of their pet. Therefore, the correct underlying algorithms must be selected that simplify external knowledge of the pet's location. Furthermore, tracking the pose of a pet is simpler than that of a human being given the typically upright posture for humans. Similarly, pets get easily occluded by furniture items in an indoor setting. Therefore, obtaining accurate ranges can become challenging.

Can more than one pet be tracked? It is possible to modify the PetTrack ranging protocol to accommodate a few pet-wearable UWB devices. However, this solution cannot support a large number of pets, such as in a pet-daycare facility. A TDoA version of the protocol will be required for supporting a very large number of pets.

Can PetTrack be used to track activities like eating, drinking? Can it be used to measure calories burned? Using the accelerometer data and the know weight and breed of the pet, we can determine rate of calorie burn. Right now, our model is pet-dependent, which means the model has to be trained to detect different types of activity for every pet that uses PetTrack. It is possible to use more domain knowledge of different pets and more advanced machine learning techniques to make this into a pet independent model in the future.

Can we improve the pose inference accuracy with other sensors? It's possible to fuse accelerometer data with other sensor inputs, such as gyroscope, UWB, light sensor, etc. to improve the inference accuracy. For example, we observed that for a dog laying on its belly and standing on its legs, the accelerometer data is very difficult to distinguish; but with a light sensor under the collar, it will help differentiate whether the dog is touching the floor or not.

Can PetTrack be used to teach pets and reinforce what their owners teach them? It is possible to define some indoor spaces as being off-limits for the pet. If, for example, the pet attempts to enter into a designated off-limits space, an ultrasonic sound can be emitted, alerting the pet of the off-limits rule and discouraging the pet from entering certain rooms, even rooms that are irregularly shaped.

8 CONCLUDING REMARKS

PetTrack is a pet tracking system that allows the pet owners to monitor pets' activities through real-time location and pose tracking. PetTrack demonstrated reliable localization and pose inference performance, which we believe will enable a wide spectrum of pet-centered applications.

REFERENCES

[1] American pet products association releases newest edition of national pet owners survey. In *American Pet Products Association Releases Newest Edition of National Pet Owners Survey.*

[2] Rita Brugarolas, David Roberts, Barbara Sherman, and Alper Bozkurt. Machine learning based posture estimation for a wireless canine machine interface. In *2013 IEEE Topical Conference on Biomedical Wireless Technologies, Networks, and Sensing Systems*, pages 10–12, 2013.

[3] 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*, 11:1549, 5 2021.

[4] Ashutosh Dhekne, Ayon Chakraborty, Karthikeyan Sundaresan, and Sampath Rangarajan. Trackio: tracking first responders inside-out. In *16th {USENIX} Symposium on Networked Systems Design and Implementation ({NSDI} 19)*, pages 751–764, 2019.

[5] Jose A. Gutierrez, Edgar H. Callaway, and Raymond L. Barrett. *IEEE 802.15.4 Low-Rate Wireless Personal Area Networks: Enabling Wireless Sensor Networks*. IEEE, 2003.

[6] M. Hazas and A. Hopper. Broadband ultrasonic location systems for improved indoor positioning. *IEEE Transactions on Mobile Computing*, 5(5):536–547, 2006.

[7] IEEE802.15.4z. IEEE standard for low-rate wireless networks–amendment 1: Enhanced ultra wideband (uwb) physical layers (phys) and associated ranging techniques. *IEEE Std 802.15.4z-2020 (Amendment to IEEE Std 802.15.4-2020)*, pages 1–174, 2020.

[8] Anton Ledergerber, Michael Hamer, and Raffaello D'Andrea. A robot self-localization system using one-way ultra-wideband communication. In *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 3131–3137. IEEE, 2015.

[9] Locatify. What is the new apple u1 chip, and why is it important? https://locatify.com/blog/what-is-the-new-apple-u1-chip-and-why-is-it-important/, 2020.

[10] Jorge J. Moré. The levenberg-marquardt algorithm: Implementation and theory. In G. A. Watson, editor, *Numerical Analysis*, pages 105–116, Berlin, Heidelberg, 1978. Springer Berlin Heidelberg.

[11] Dries Neirynck, Eric Luk, and Michael McLaughlin. An alternative double-sided two-way ranging method. In *2016 13th workshop on positioning, navigation and communications (WPNC)*, pages 1–4. IEEE, 2016.

[12] Michael Stocker, Bernhard Großwindhager, Carlo Alberto Boano, and Kay Römer. SnapLoc: an ultra-fast UWB-based indoor localization system for an unlimited number of tags. In *Proceedings of the 18th International Conference on Information Processing in Sensor Networks*, pages 348–349, New York, NY, USA, apr 2019. ACM.

[13] Songsri Tangsripairoj, Parit Kittirattanaviwat, Kamonwan Koophiran, and Lalita Raksaithong. Bokk meow: A mobile application for finding and tracking pets. In *2018 15th International Joint Conference on Computer Science and Software Engineering (JCSSE)*, pages 1–6, 2018.

[14] 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.

[15] Deepak Vasisht, Swarun Kumar, and Dina Katabi. Decimeter-level localization with a single wifi access point. In *Proceedings of the 13th Usenix Conference on Networked Systems Design and Implementation*, NSDI'16, page 165–178, USA, 2016. USENIX Association.