



OPEN

DATA DESCRIPTOR

# CAPTURE-24: A large dataset of wrist-worn activity tracker data collected in the wild for human activity recognition

Shing Chan<sup>1,2</sup>, Yuan Hang<sup>1,2</sup>, Catherine Tong<sup>1,2,3</sup>, Aidan Acquah<sup>1,2,4</sup>, Abram Schonfeldt<sup>1,2</sup>, Jonathan Gershuny<sup>5</sup> & Aiden Doherty<sup>1,2</sup>✉

Existing activity tracker datasets for human activity recognition are typically obtained by having participants perform predefined activities in an enclosed environment under supervision. This results in small datasets with a limited number of activities and heterogeneity, lacking the mixed and nuanced movements normally found in free-living scenarios. As such, models trained on laboratory-style datasets may not generalise out of sample. To address this problem, we introduce a new dataset involving wrist-worn accelerometers, wearable cameras, and sleep diaries, enabling data collection for over 24 hours in a free-living setting. The result is CAPTURE-24, a large activity tracker dataset collected in the wild from 151 participants, amounting to 3883 hours of accelerometer data, of which 2562 hours are annotated. CAPTURE-24 is two to three orders of magnitude larger than existing publicly available datasets, which is critical to developing accurate human activity recognition models.

## Background & Summary

With the increasing adoption of activity trackers such as Fitbit and Apple Watch, the ability to extract objective health-related behavioural insights at an unprecedented scale prompts new opportunities in medicine. A particularly promising direction is the use of accelerometer-based activity recognition in healthcare, where it is still common to use recall diaries or time and labour-intensive methods such as in-person clinical fitness tests. These approaches suffer from subjectivity and/or scalability issues. Wrist-worn accelerometers, being low-cost, low-powered and convenient, allow us to efficiently obtain an objective and high-resolution picture of a user's daily activities, enabling new opportunities for real-time precision medicine, digital phenotyping for routine care and clinical trials<sup>1-3</sup>, and large-scale population and epidemiological studies<sup>4-6</sup>. The success of these applications depends on a reliable activity recognition model, requiring a sizeable and representative labelled dataset for model development.

However, existing open accelerometer datasets have many shortcomings due to the commonly used data collection protocol, whereby participants are brought into a controlled environment to perform a series of pre-defined activities in a specified order and under supervision. This lab-style setup leads to the following issues: 1) the data collected is usually limited due to the labour-intensive nature of the approach; 2) it often excludes mixed activities; 3) even when mixed activities are included, the study's design (prompted instructions, supervised performance, controlled environment) encourages uniform and stereotypical movement patterns due to participants' tendency to conform (acquiescence bias); 4) the sequence of activities is unnatural. The latter means that such dataset cannot be used for sequence modelling (e.g. hidden Markov models, recurrent neural networks) – for instance, if the study were designed with a fixed sequence of activities (e.g. “office work” → “standing” → “walking” → “stairs”), sequence modelling would be overfitted to this specific order (the transition matrix would be sparse and close to identity), showing high in-dataset test accuracy but failing to generalise outside the dataset. Note that randomising the sequence would eliminate any temporal dependencies,

<sup>1</sup>Big Data Institute, University of Oxford, Oxford, UK. <sup>2</sup>Nuffield Department of Population Health, University of Oxford, Oxford, UK. <sup>3</sup>Department of Computer Science, University of Oxford, Oxford, UK. <sup>4</sup>Department of Engineering Science, University of Oxford, Oxford, UK. <sup>5</sup>Social Research Institute, University College London, London, UK. ✉e-mail: [aiden.doherty@ndph.ox.ac.uk](mailto:aiden.doherty@ndph.ox.ac.uk)

thereby undermining the purpose of sequence modelling. We emphasise this point because human activities are inherently sequential, and modelling temporal dependencies is crucial for achieving accurate classifiers, as evidenced by both the literature and the model benchmarks presented in this work.

The overall impact of the aforementioned limitations in training machine learning models is that the resulting classifier appears to perform well within the confines of a study, even with proper validation and testing, but fails to deliver the same performance in the real world. The root of this problem lies in the unrepresentativeness of dataset. To address this, we release CAPTURE-24 – a large, in-the-wild dataset of annotated wrist-worn accelerometer readings from 151 participants over approximately 24 hours, amounting to 2562 hours of annotated activity data which is many times larger than existing datasets (Table 1). The dataset was initially created to compare time-use diaries with device measurements<sup>7</sup>. Thus, it prioritizes data collection under unscripted free-living settings and routine-capturing durations (i.e., at least 24 hours).

We anticipate the CAPTURE-24 dataset to be a valuable resource in wearable sensor-based human activity recognition, especially for research in data-hungry methods such as deep learning. We illustrate this in our benchmarks, which include commonly used methods such as random forest, XGBoost, hidden Markov models, and deep learning methods.

**Comparison to Existing Datasets.** Table 1 illustrates how CAPTURE-24 compares to existing publicly available wrist-worn accelerometer datasets:

- **Dataset Size and Duration:** CAPTURE-24 contains 2562 hours of annotated data collected from 151 participants, each monitored for approximately 24 hours. This is substantially larger than other datasets, typically containing less than 50 hours of data collected from much fewer participants.
- **Annotation Granularity and Diversity:** The CAPTURE-24 dataset provides over 200 fine-grained activity annotations, covering a wide range of leisure, sports, and occupational activities. Only Opportunity++ dataset has a higher number of annotations.
- **Free-Living Data Collection:** Unlike many datasets that collect data in controlled environments, CAPTURE-24 was collected in a free-living setting. This means participants carried out their normal daily activities, providing a more realistic and heterogeneous dataset. This aspect is crucial as it ensures that models trained on CAPTURE-24 can better generalise to real-world applications.

*Large but Unlabelled Datasets.* In addition to the above datasets, there are several large-scale, unlabelled datasets with consumer-grade<sup>4,8</sup> and research grade devices<sup>9,10</sup> that contribute significantly to population health studies and wearable sensor research. We focus on research-grade devices that offer raw accelerometer data.

- **UK Biobank:** This is a major resource for health research, containing extensive data from over 500,000 participants. While the focus is on a wide range of health-related information, it includes large-scale unlabelled accelerometer data that can be valuable for training and validating activity recognition models.
- **NHANES (National Health and Nutrition Examination Survey):** Conducted by the Centers for Disease Control and Prevention (CDC) in the United States, NHANES collects health and nutritional data from thousands of participants, including unlabelled accelerometer data.
- **China-Kadoorie Biobank:** This dataset encompasses health data from over 500,000 Chinese participants. Similar to the UK Biobank, it includes unlabelled accelerometer data that provides a vast resource for developing and testing new activity recognition models.

These large unlabelled datasets, while not providing direct annotations, offer a wealth of raw data that can be leveraged through semi-supervised or unsupervised learning techniques to enhance the performance of activity recognition systems. They also provide a broad demographic and geographic diversity, enabling more generalized and robust model development.

## Methods

**Data Acquisition.** The CAPTURE-24 study was the first sizeable attempt to test traditional self-report time-use diaries against real-time passive sensing instruments, namely, wearable cameras and activity trackers<sup>7</sup>. Data collected from this study (carried out in 2014–2015) forms the majority of our CAPTURE-24 dataset – the data being released includes additional data collected since then. Additional processing, labelling for activity recognition, and anonymization were conducted to permit its open release. An overview of the procedures carried out is depicted in Fig. 1.

The design and associated operating procedures of the CAPTURE-24 study were based on the findings of a pilot study ( $n = 14$ )<sup>11</sup>. The ethical approval was granted by the University of Oxford Inter-Divisional Research Ethics Committee (Ref SSD/CUREC1A/13–262). Members of the public from Oxfordshire, United Kingdom, were recruited as study participants following advertisements with a £20 voucher for taking part. A member of the research team met with participants to explain the project purpose, gain written informed consent, complete a short demographic questionnaire (including gender, age, height and weight) and deliver the instruments. During the designated data collection day, participants were asked to wear a wrist-worn accelerometer continuously and a wearable camera while awake. For sleep monitoring, participants were asked to complete a simple sleep diary consisting of two questions: “what time did you first fall asleep last night?” and “what time did you wake up today (eyes open, ready to get up)?”. Participants were also asked to keep a harmonised European time-use survey<sup>12</sup>, from which sleep information was extracted when data was missing from the simple sleep diary.

Dataset	Size	Annotations†	Domain	Free-living	Year
CAPTURE-24	2562 hrs (151 ppl × 24 hrs) <sup>‡</sup>	>200	Leisure, Sports, Occupation	✓	2024
WISDM <sup>48</sup>	43 hrs (51 ppl × 54 mins)	18	Leisure	✗	2019
REALDISP <sup>49</sup>	39.8 hrs (17 ppl × 140.5 mins)	33	Sports	✗	2012
OxWalk <sup>50</sup>	39 hrs (39 ppl × 60 mins)	Steps	Leisure	✓	2023
Hang-Time <sup>51</sup>	37.8 hrs (24 ppl × 94.4 mins)	11	Sports	✓	2023
Leisure Activities <sup>52</sup>	27.7 hrs (6 ppl × 277 mins)	6	Leisure	✓	2012
Clemson <sup>53</sup>	22.5 hrs (30 ppl × 45 mins)	Steps	Leisure	✗	2017
WetLab <sup>54</sup>	21.0 hrs (22 ppl × 57.2 mins)	8	Occupation	✗	2015
Realworld <sup>55</sup>	18 hrs (15 ppl × 70 mins)	8	Leisure	✗	2016
Smartwatch Swimming <sup>56</sup>	17 hrs (40 ppl × 25.5 mins)	5	Sports	✗	2019
WEAR <sup>57</sup>	15 hrs (18 ppl × 50 mins)	18	Sports	✓	2023
TNDA-HAR <sup>58</sup>	5.7 hrs (23 ppl × 14.6 mins)	8	Leisure	✗	2021
Opportunity++ <sup>59</sup>	5.3 hrs (4 ppl × 80 mins)	>24,000	Leisure, Gestures	✗	2021
Opportunity <sup>60</sup>	5.3 hrs (4 ppl × 80 mins)	13	Leisure, Gestures	✗	2010
UC Berkley WARD <sup>61</sup>	5 hrs (20 ppl × 15 mins)	13	Leisure	✗	2009
PAMAP2 <sup>62</sup>	4.5 hrs (9 ppl × 30 mins)	18	Leisure	✗	2012
CMU-MMAC <sup>63</sup>	3.6 hrs (43 ppl × 5 mins)	5	Cooking	✗	2008
Skoda <sup>64</sup>	3 hrs (1 ppl × 180 mins)	10	Occupation	✗	2008
ADL <sup>65</sup>	2.7 hrs (16 ppl × 10 mins)	14	Leisure	✗	2014
MHEALTH <sup>66</sup>	2.5 hrs (10 ppl × 15 mins)	12	Leisure	✗	2014
UTD-MHAD <sup>67</sup>	1.4 hrs (8 ppl × 10.8 mins)	27	Gestures	✗	2015
UC Berkley HMAD <sup>68</sup>	1.3 hrs (12 ppl × 6.3 mins)	11	Leisure, Gestures	✗	2013
Daily and Sports Activities <sup>69</sup>	0.7 hrs (8 ppl × 5 mins)	19	Leisure, Sports	✗	2010
UTD-MHAD KinectV2 <sup>70</sup>	0.4 hrs (6 ppl × 4.3 mins)	10	Gestures	✗	2015

**Table 1.** Publicly available wrist-worn accelerometer datasets. hrs: hours; ppl: people; mins: minutes. †For datasets with annotations at different levels of granularity, the largest number of annotations at a given level of granularity was chosen. ‡We released 3883 hours of total recording, of which 2562 hours are labelled.

An initial 166 participants were recruited, of which 151 remained after disregarding participants with incomplete, corrupted and/or bad quality data.

**Accelerometer.** Participants were asked to wear an Axivity AX3 wrist-worn tri-axial accelerometer on their dominant hand. The accelerometer was set to capture tri-axial acceleration data at 100 Hz with a dynamic range of  $\pm 8g$ . Axivity device has been validated for estimating energy expenditure in a free-living environment<sup>13</sup>. This device has also demonstrated equivalent signal vector magnitude output on multi-axis shaking tests with other commonly-used accelerometers<sup>14</sup>.

**Wearable Camera.** Wearable cameras were used to collect ground truths of the participants' activities while wearing the accelerometers. Participants were given an OMG Life Autographer, a wearable camera worn around the neck which automatically takes photographs every 20 - 40 seconds and has up to 16 hours battery life and storage capacity for over one week's worth of images<sup>15</sup>. When worn, the camera is reasonably close to the wearer's eye line and has a wide-angle lens to CAPTURE the wearer's view<sup>16</sup>. Previously, annotations from wearable cameras have been found to have strong agreement with the more expensive direct observation methods to classify activity types (inter-rater reliability via Cohen's  $\kappa$  of 0.92)<sup>11</sup>. Recently, the camera images have also been found to correctly identify 85% + of sitting time against direct observations<sup>17</sup>. Some sample images captured by the wearable camera can be seen in Fig. 1.

Due to the intrusive nature of wearable cameras, we abide by the ethical framework established by Kelly and colleagues through data collection and processing; this included scheduling a reviewing session with participants who revisited their own camera data to remove any unwanted or sensitive images<sup>18</sup>. The public CAPTURE-24 dataset also excludes image data from the wearable cameras – only text annotations of the images are provided.

**Data Annotation.** We relied on the time-stamped wearable camera images to annotate the accelerometer data during wake time, and sleep diaries during sleep time. To standardize the annotation taxonomy, we employed activity codes from the Compendium of Physical Activities (CPA)<sup>19</sup>. This describes activities and their contexts in a hierarchical fashion with an associated Metabolic Equivalent of Task (MET) score to represent the mass-specific energy expenditure of activities. An example found in CAPTURE-24 is "occupation; interruption; 11795 walking on job and carrying light objects such as boxes or pushing trolleys; MET 3.5". To ensure the reliability of the annotation process, all annotators had to complete a short training course. This covered ethics training for handling image data, usage of a specifically-developed

image browsing software<sup>20</sup>, annotation training, and finally passing annotation quality checks on a held-out gold-standard dataset, where annotators have to achieve an (Cohen's) inter-rater agreement score of  $\kappa > 0.8$ .

**Data Processing.** *Data Extraction.* The Axivity Omgui software (<https://github.com/digitalinteraction/openmovement/wiki/AX3-GUI>) distributed by the accelerometer manufacturers was used for initialization of the measurements, synchronization, and downloading binary accelerometry files recorded on the devices. On top of this, we applied sampling rate correction by nearest-neighbor interpolation to fix any irregular sampling that may happen due to machine error. To correct for any accelerometer miscalibration and measurement drifts, gravity autocalibration<sup>21</sup> was applied to reduce discrepancies across devices.

*De-identification.* To protect participant privacy, we selected a subset of the collected data for public release – the accelerometer data, the text annotations, and the participants' gender and age. Image data is not included in the release. Participant ages were binned into 4 similarly-sized groups {"18–29", "30–37", "38–52" and "53 or above"}. For further de-identification, actual dates were randomized and timestamps were shifted by a small random amount. We also reconsidered the acquired CPA code annotations containing sensitive or rare activities, judging on a case-by-case basis whether to simplify the annotation to ensure anonymization. A hypothetical example is a code containing the description "skiing, cross country, >8.0 mph, elite skier, racing; MET 15" which comes from only one participant in the dataset (a professional skier), the annotation will be re-labeled as "sports; MET 15". The inclusion of the original label's MET ensures that the activity intensity is still retained in the annotation, but only in cases where its inclusion does not permit the activity to be uniquely identified independently from the CPA, in which case the MET value is rounded to the nearest, most commonly occurring MET value seen in CAPTURE-24.

**Benchmarks.** We considered two activity recognition tasks: one to classify intensity levels of physical activity, and another to classify activities of daily living. For these, we re-worked the CPA codes accordingly, mapping them to two sets of simplified labels<sup>5,6</sup>. The labels for intensities of physical activity are {"sleep", "sedentary", "light physical activity", "moderate-to-vigorous physical activity"}, and the labels for activities of daily living are {"sleep", "sitting", "standing", "household-chores", "manual-work", "walking", "mixed-activity", "vehicle", "sports", "bicycling"}.

*Data preprocessing for activity recognition.* We followed the common practice of sliding-window segmentation<sup>22</sup> to extract fixed-size, non-overlapping windows of ten seconds. This resulted in a dataset of  $n = 922,199$  windows in total, each with dimension (3, 1000) (10 sec, 100 Hz, tri-axial). Data from 100 participants (P001–P100, with 618, 129 windows) were used for model derivation, while the rest (P101–P151, with 304, 070 windows) was set aside for model evaluation; we refer to these as the Derivation Set and Test Set respectively. The class distribution in both sets remained similar (Fig. S2).

*Models.* We consider the following commonly used methods in the activity recognition literature:

- **Random forest (RF)** We used a balanced random forest<sup>23</sup> with 3000 trees. The number of trees were chosen to be as large as possible. The model was very robust to the remaining hyperparameter choices, so we used the recommended default values.
- **XGBoost** We used XGBoost<sup>24</sup> and Bayesian optimization<sup>25</sup> to tune the hyperparameters (number of estimators, max depth, gamma, regularization coefficients) with 100 iterations, although we found that it did not significantly improve upon the default hyperparameters.
- **Convolutional neural network (CNN)** We use 1D convolutions, residual blocks<sup>26</sup>, and anti-aliased down-sampling<sup>27</sup>. We tuned the kernel sizes, number of blocks, and number of filters using grid search and the ASHA scheduler<sup>28</sup>. See Supplementary Section B for further details.
- **Recurrent neural network (RNN)** The backbone of the architecture is the CNN mentioned above, with the second last layer (originally a fully-connected layer) replaced by a bidirectional Long Short-Term Memory module<sup>29</sup>. This model can take a sequence of windows to model temporal dependencies. A maximum sequence length of 8 (80 sec) is used. See Supplementary Section B for further details.
- **Hidden Markov models (HMM)** Additionally, we consider the application of hidden Markov models on top of all aforementioned models to model the temporal dependencies between windows. The HMM is applied post-hoc to the final sequence of outputs from the base models. Note that while RNN is already a temporal model, we found further improvements when applying HMM on top.

For RF and XGBoost, we extracted time-series features from the accelerometry that are commonly used in the literature<sup>30</sup> including time and frequency domain features, angular and peak features, resulting in a total of 40 features per window. See Supplementary Section A for the full list of features.

*Metrics.* The distribution of activities was highly imbalanced reflecting the free-living nature of the collected data ("sleeping", "sitting" and "standing" make up more than 60% of activities). As a result, we reported our evaluations using metrics often deemed more intuitive for this scenario such as macro-averaged F1-score, Cohen's  $\kappa$ , and Pearson-Yule's  $\phi$  coefficient<sup>31,32</sup> (also known as the Matthews correlation coefficient) on the test set. We use bootstrapping ( $n = 100$ ) to estimate 95% confidence intervals<sup>33</sup> on all reported metrics.

*Training details.* In the deep learning experiments, we further split the derivation set of 100 users into 80 users (503, 880 windows) for training and 20 users (125, 970 windows) for validation and early stopping. A batch



**Fig. 1** Overview of the creation of the CAPTURE-24 Dataset. Recruited subjects wore an activity tracker for roughly 24 hours. They also wore a camera during daytime and used a diary to register their sleep times during nighttime. The collected data was processed and harmonised to obtain acceleration time-series data annotated with the activities performed. CPA: Compendium of Physical Activities; MET: metabolic equivalent. Note that the camera images are not part of the dataset release.

size of 512 was used throughout except for the RNN model where it was reduced to 64 in response to increased computational burden due to the sequence length of 8. Stochastic gradient descent with restarts<sup>34,35</sup> was used for optimization. Four data augmentation methods were explored<sup>36</sup>: jittering, time warping, magnitude warping, and shifting. See Supplementary Section B for further details.

### Data Records

Our dataset is hosted at the Oxford University Research Archive under the Creative Commons “Attribution 4.0 International (CC BY 4.0)” License, at <https://doi.org/10.5287/bodleian:NGx0JOMP537>. The raw accelerometry data has been processed and stored as compressed CSV files using the `biobankAccelerometerAnalysis` tool. For each participant, the raw accelerometry file contains the following columns:

- Time: the timestamp for each accelerometry reading in milliseconds;
- X, Y, Z: the raw accelerometry along each of the axes in g;
- Annotation: the activity annotation using a category from the Compendium of Physical Activities.

In addition, we also provided the “annotation-label-dictionary.CSV” to provide the annotation mapping from fine-grained activity to high-level classes that be for machine learning, genetics and population health studies<sup>5,38,39</sup>. Finally, age group and sex information for each participant are stored in “metadata.CSV”.

### Technical Validation

Initially, 166 participants were recruited. After discarding incomplete, corrupted and/or bad quality data, **151 participants** remained amounting to a total of **3883 hours** of data. The different data sources (activity tracker, camera, sleep diary) were then harmonised and processed, resulting in **2562 hours** of annotated data. Participant demographics are summarised in Table 2. The majority of participants were women (66%). Different age groups were relatively well-represented, which is important for developing models that generalize well for changes in movement patterns due to aging (e.g. walking pace).

A total of **206 unique CPA codes** were identified. The CPA codes followed a long-tail distribution (Supplementary Section C, dominated by the “sleeping” activity which constitute more than a third of activities. The most and least frequent CPA codes are shown in Table 3. As the 206 codes can be overly detailed, we devised six schema (included in the data release) for mapping the fine-grained codes into sets of simplified labels.

Each scheme has an intended use according to a research question. For example, for an epidemiological study focusing on physical activity levels, it may be convenient to summarise the codes into 4 classes: “sleep”, “sedentary activity”, “light physical activity”, “moderate-to-vigorous physical activity”. For a general activity recognition

	All	Derivation Set	Test Set
	<i>n</i> (%)	<i>n</i> (%)	<i>n</i> (%)
<b>Gender</b>			
Male	52 (34.4)	36 (36.0)	16 (31.4)
Female	99 (65.6)	64 (64.0)	35 (68.6)
<b>Age</b>			
18–29	43 (28.5)	27 (27.0)	17 (33.3)
30–37	37 (24.5)	26 (26.0)	14 (27.5)
38–52	37 (24.5)	24 (24.0)	10 (19.6)
≥53	34 (22.5)	23 (23.0)	10 (19.6)

**Table 2.** Demographic information for CAPTURE-24 participants.

study, we may instead consider activities such as “walking”, “standing”, “bicycling”. Table 4 shows three schema used in previous works to answer different research questions<sup>5,6,38</sup>.

**Benchmarks.** Results for the different models are summarised in Table 5. Scores for the classification of physical activity levels are shown in Table 5a, and those for the classification of daily-living activities are shown in Table 5b. In each subtable, top half shows scores for the base models, and bottom half shows scores for the models enhanced with HMM smoothing. We see that HMM consistently achieves big improvements across tasks, models and metrics, highlighting the importance of modelling temporal dependencies. Further, we found that RF + HMM and XGBoost + HMM are already competitive, both performing on par or better than the more expensive models without HMM. The importance of temporal modelling was also seen within the models without HMM, where RNN excelled as it had the context of up to 8 consecutive windows to make predictions. Notably, we found that HMM further improved upon RNN, suggesting that longer sequence modelling would be fruitful. Interestingly, we found CNN + HMM to be the best performing model overall even though RNN performed better in the non-HMM cases.

*Challenges in Activity Recognition in the Wild.* We found that scores for the classification of daily-living activities were consistently lower than those for the classification of physical activity levels. More granular classification is generally harder for all types of tasks, but it is especially so with our dataset due to the ambiguity of many free-living human activities. Figure 2 shows the confusion matrices using the RF + HMM model. For activities of daily living, most of the confusion happens between the activities “household-chores”, “standing”, “walking”, “manual-work” and “mixed-activity”. This is expected given that, in free-living settings, these activities are naturally intertwined (e.g. the household chore “cleaning, sweeping carpet or floor” inevitably involves some degree of “walking” and “standing”), as opposed to data collected in laboratory settings where the scripted activities tend to be clearly segmented and the movement patterns show less heterogeneity. Regarding classification of physical activity intensity levels, their definitions from real-world human activities tend to be less ambiguous, therefore we observed less confusion for this classification task.

*Performance against Dataset Size.* We highlighted the importance of having large datasets for data-intensive deep learning methods. We assessed the performance as a function of dataset size by running the benchmarks on varying number of subjects included in the derivation set (the test set of 51 subjects is unchanged). From Fig. 3, we observed that in the small-data regime the outperformance of deep learning models is not so clear. In particular, if we consider only non-temporal models (RF, XGBoost and CNN), we could see that CNN only starts to outperform after around 40 subjects ( $\approx 650$  person-hours). Similarly for the temporal models (RNN and \*-HMM), the clear advantage of CNN-HMM became apparent only after around 30 subjects ( $\approx 500$  person-hours).

## Usage Notes

We presented the CAPTURE-24 dataset to address shortcomings of existing activity recognition datasets – namely, limited dataset sizes and unrepresentativeness due to intrusive data collection methods resulting in short time spans, limited and scripted activities, low pattern heterogeneity, and manufactured activity sequences. We described in detail how CAPTURE-24 addressed these issues with a novel collection protocol involving indirect measurements using wearable cameras and sleep diaries, allowing for long time spans (24 hours or more) in real-world settings while also being less labour-intensive and more scalable. We also described procedures taken to comply with privacy and ethics standards to permit the public release of the dataset. With 2562 hours of annotated data (and 3883 hours overall), the released dataset is 2 to 3 orders of magnitude larger than existing public accelerometer datasets.

We presented benchmarks for activity recognition on this dataset with commonly used methods in the literature and discussed challenges for activity recognition in the wild. In particular, we highlighted challenges in activity recognition in the wild as many activities in the real world are intertwined, in contrast to those collected in laboratory settings. We also highlighted the importance of having large HAR datasets for deep learning research, suggesting that existing dataset sizes are insufficient to achieve the full potential of their methods, rendering any model comparison unreliable.

**Limitations.** The CAPTURE-24 only contains a convenience sample of participants in Oxford. Therefore, larger datasets using more diverse populations are needed. For example, a similar dataset was collected in China

Most frequent
7030 sleeping;MET 0.95
occupation;office and administrative support;11580 office/computer work general;MET 1.5
home activity;household chores;preparing meals/cooking/washing dishes;5035 kitchen activity;MET 3.3
home activity;miscellaneous;sitting;9060 sitting/lying reading without observable activities;MET 1.3
occupation;office and administrative support;11580 office wok/computer work general;MET 1.5
Least frequent
home activity;household chores;5140 sweeping garage sidewalk or outside of house;MET 4.0
carrying heavy loads;MET 8.0
occupation;interruption;walking;17070 walking downstairs;MET 3.5
leisure;miscellaneous;walking;17070 descending stairs;MET 3.5
occupation;interruption;walking;17133 walking upstairs;MET 4.0
leisure;miscellaneous;21070 (generic) walking/standing combination indoor;MET 3.0

**Table 3.** Annotation examples.

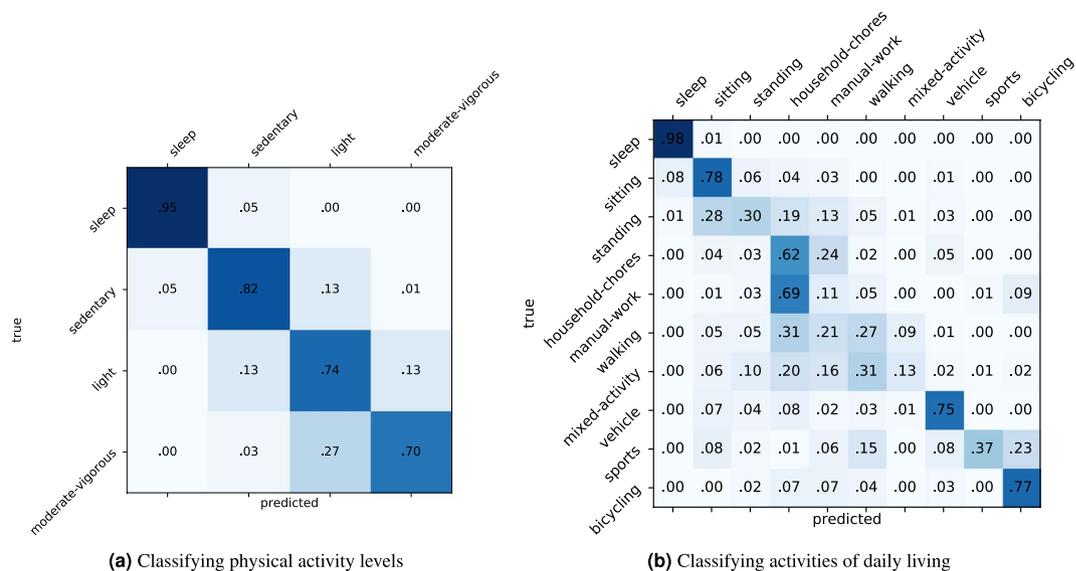
Reference	Study objective	Labels
Willets <i>et al.</i> <sup>5</sup>	Activity recognition	bicycling, sports, vehicle, mixed, walking, manual-work, household-chores, standing, sitting, sleep
Willets <i>et al.</i> <sup>5</sup>	Activity recognition	bicycling, sit/stand, walking, vehicle, mixed activity, sleep
Doherty <i>et al.</i> <sup>38</sup>	Genomic discovery	sleep, sedentary, walking, moderate physical activity
Walmsley <i>et al.</i> <sup>6</sup>	Epidemiology	sleep, sedentary behaviour, light physical activity, moderate-to-vigorous physical activity

**Table 4.** Labels for classification tasks on CAPTURE-24 used in previous studies and their intended objectives.

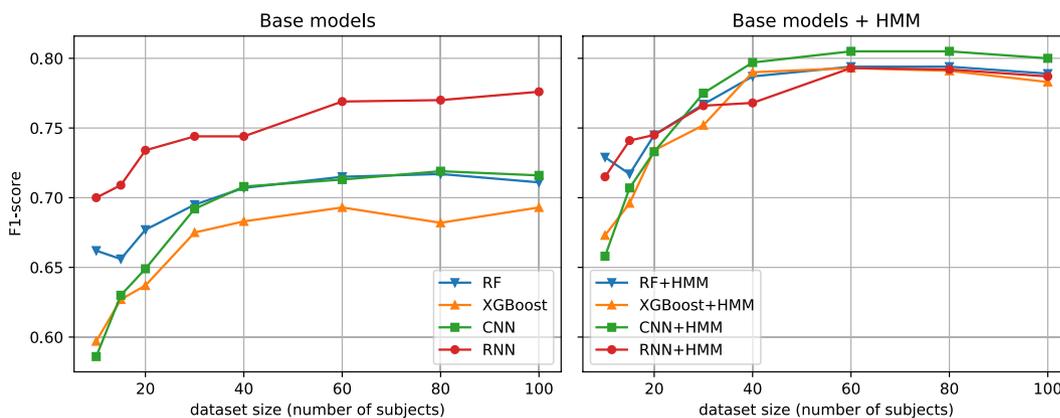
(a) Classifying physical activity levels			
Model	F1-score	Cohen's $\kappa$	Pearson-Yule's $\phi$
RF	.711 (.710, .713)	.665 (.664, .667)	.666 (.664, .668)
XGBoost	.693 (.691, .695)	.666 (.664, .668)	.667 (.665, .669)
CNN	.716 (.714, .718)	.704 (.703, .706)	.704 (.702, .705)
RNN	.776 (.774, .778)	.790 (.788, .791)	.790 (.789, .792)
RF + HMM	.789 (.788, .791)	.772 (.771, .774)	.773 (.771, .774)
XGBoost + HMM	.783 (.781, .785)	.772 (.771, .774)	.773 (.771, .774)
CNN + HMM	.800 (.798, .802)	.814 (.812, .816)	.815 (.813, .816)
RNN + HMM	.787 (.785, .789)	.800 (.798, .801)	.801 (.799, .802)
(b) Classifying activities of daily living			
Model	F1-score	Cohen's $\kappa$	Pearson-Yule's $\phi$
RF	.388 (.386, .390)	.514 (.512, .516)	.519 (.518, .521)
XGBoost	.399 (.396, .401)	.576 (.574, .577)	.580 (.579, .582)
CNN	.485 (.482, .488)	.626 (.624, .627)	.627 (.625, .629)
RNN	.534 (.530, .536)	.697 (.695, .698)	.698 (.696, .700)
RF + HMM	.486 (.484, .489)	.637 (.636, .639)	.641 (.639, .643)
XGBoost + HMM	.483 (.480, .486)	.682 (.680, .684)	.686 (.684, .688)
CNN + HMM	.576 (.574, .580)	.735 (.733, .736)	.737 (.736, .739)
RNN + HMM	.537 (.535, .540)	.714 (.713, .716)	.716 (.715, .718)

**Table 5.** Performance for activity recognition on the CAPTURE-24 Dataset. 95% confidence intervals are shown in brackets.

for human activity recognition as part of the China Kadoorie Biobank wearable study<sup>9</sup>. As wearable sensing technologies improve, multi-modal monitoring for human activity recognition over time has become feasible, improving the predictive power for labour-intensive activity with little wrist movement and the classification of the sleep stages.



**Fig. 2** Confusion matrix for random forest + hidden Markov model: **(a)** performance on high-level activities **(b)** performance on activities with more granularity.



**Fig. 3** F1-score as a function of dataset size for physical activity classification: **(a)** base models **(b)** base models + Hidden Markov Models.

Furthermore, camera data may sometimes be uninformative for annotation due to obstruction, poor lighting conditions and blurriness. Since the cameras record data at a low frame rate ( $\approx 0.03\text{Hz}$ ) – much lower than that of the accelerometers ( $100\text{Hz}$ ) – activities could have been missed. As a result, it is possible that the annotators may assign CPA codes through guess work despite our best efforts in covering uncertain scenarios in the annotator training. A further limitation with CPA codes is that they were originally developed for use in epidemiological studies to standardise the assignment of MET values in physical activity questionnaires, thus some codes place more emphasis on distinguishing energy intensities rather than behaviours. This results in some CPA codes being ambiguous for retrospective interpreting and re-labelling. For example, the code "home activity; miscellaneous; standing; 9050 standing talking in person/ on the phone/ computer (skype chatting) or using a mobile phone/ smartphone/ tablet; MET 1.8" precludes distinguishing whether the participant was speaking to someone in person or through a specific device, which might have been useful in studies looking to understand people's screen-time or social behaviours. Our existing benchmark incorporates common methods used for HAR, future work could also benefit from leveraging more recent modeling techniques using DeepConvLSTM<sup>40,41</sup>, transformers<sup>42</sup>, and self-supervised learning<sup>43–46</sup>.

**Research directions.** Although the feasibility of activity recognition solely from accelerometer data has been debated in recent work<sup>47</sup>, a proper investigation has been lacking due to the lack of realistic datasets. As mentioned, current datasets have limited and homogeneous activities which optimistically bias any assessment. The CAPTURE-24 could be useful for such investigation. The problem of open set recognition is another interesting direction in HAR as there are practically infinite number of activities that one could consider. The fine-grained and hierarchical annotations in CAPTURE-24 can be leveraged to study this problem within the

framework of zero-shot and few-shot learning. Finally, we saw that temporal modelling using hidden Markov models consistently improved performance, but other time-series methods<sup>40,42</sup> could be investigated leveraging the unique long-span aspect of our dataset.

### Code availability

All the code associated with this project can be used free of charge for non-commercial purposes. <https://github.com/OxWearables/oxford-wearable-camera-browser> holds the code for the annotation tool. The data analytics script can be accessed via biobankAccelerometerAnalysis repository: <https://github.com/OxWearables/biobankAccelerometerAnalysis>. Finally, the machine learning benchmarks can be accessed via: <https://github.com/activityMonitoring/capture24>.

Received: 1 March 2024; Accepted: 2 October 2024;

Published online: 16 October 2024

### References

- Creagh, A. P. *et al.* Digital health technologies and machine learning augment patient reported outcomes to remotely characterise rheumatoid arthritis. *MedRxiv* 2022–11 (2022).
- Schalkamp, A.-K., Peall, K. J., Harrison, N. A. & Sandor, C. Wearable movement-tracking data identify parkinson's disease years before clinical diagnosis. *Nature Medicine* 1–9 (2023).
- Gupta, A. S., Patel, S., Premasiri, A. & Vieira, F. At-home wearables and machine learning sensitively capture disease progression in amyotrophic lateral sclerosis. *Nature Communications* 14, 5080 (2023).
- Master, H. *et al.* Association of step counts over time with the risk of chronic disease in the all of us research program. *Nature medicine* 28, 2301–2308 (2022).
- Willets, M., Hollowell, S., Aslett, L., Holmes, C. & Doherty, A. Statistical machine learning of sleep and physical activity phenotypes from sensor data in 96,220 uk biobank participants. *Scientific reports* 8, 1–10 (2018).
- Walmsley, R. *et al.* Reallocating time from machine-learned sleep, sedentary behaviour or light physical activity to moderate-to-vigorous physical activity is associated with lower cardiovascular disease risk. *medRxiv* (2020).
- Gershuny, J. *et al.* Testing self-report time-use diaries against objective instruments in real time. *Sociological Methodology* 50, 318–349 (2020).
- Jonasdottir, S. S., Bagrow, J. & Lehmann, S. Sleep during travel balances individual sleep needs. *Nature Human Behaviour* 6, 691–699 (2022).
- Chen, Y. *et al.* Device-measured movement behaviours in over 20,000 china kadoorie biobank participants. *International Journal of Behavioral Nutrition and Physical Activity* 20, 138 (2023).
- Doherty, A. *et al.* Large scale population assessment of physical activity using wrist worn accelerometers: The uk biobank study. *PLoS one* 12, e0169649 (2017).
- Kelly, P. *et al.* Developing a method to test the validity of 24 hour time use diaries using wearable cameras: a feasibility pilot. *PLoS One* 10, e0142198 (2015).
- of the European Commission, S. O. *et al.* Harmonised european time use surveys, 2008 guidelines. Office for Official Publications of the European Communities (2009).
- White, T. *et al.* Estimating energy expenditure from wrist and thigh accelerometry in free-living adults: a doubly labelled water study. *International journal of obesity* 43, 2333–2342 (2019).
- Ladha, C., Ladha, K., Jackson, D. & Olivier, P. Shaker table validation of openmovement ax3 accelerometer. In *Ahmerst (ICAMPAM 2013 AMHERST): In 3rd International Conference on Ambulatory Monitoring of Physical Activity and Movement*, 69–70 (2013).
- Doherty, A. R. *et al.* Wearable cameras in health: the state of the art and future possibilities. *American journal of preventive medicine* 44, 320–323 (2013).
- Hodges, S. *et al.* Sensecam: A retrospective memory aid. In *International Conference on Ubiquitous Computing*, 177–193 (Springer, 2006).
- Martinez, J. *et al.* Validation of wearable camera still images to assess posture in free-living conditions. *Journal for the measurement of physical behaviour* 4, 47–52 (2021).
- Kelly, P. *et al.* Ethics of using wearable camera devices in health behaviour research. *Am J Prev Med* 44, 314–319 (2013).
- Ainsworth, B. E. *et al.* 2011 compendium of physical activities: a second update of codes and met values. *Med Sci Sports Exerc* 43, 1575–1581 (2011).
- Doherty, A. R., Moulin, C. J. & Smeaton, A. F. Automatically assisting human memory: A sensecam browser. *Memory* 19, 785–795 (2011).
- Van Hees, V. T. *et al.* Autocalibration of accelerometer data for free-living physical activity assessment using local gravity and temperature: an evaluation on four continents. *Journal of applied physiology* 117, 738–744 (2014).
- Bulling, A., Blanke, U. & Schiele, B. A tutorial on human activity recognition using body-worn inertial sensors. *ACM Computing Surveys (CSUR)* 46, 1–33 (2014).
- Chen, C., Liaw, A. & Breiman, L. Using random forest to learn imbalanced data. Tech. Rep. 666, University of California, Berkeley (2004).
- Chen, T. & Guestrin, C. Xgboost: A scalable tree boosting system. In *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining*, 785–794 (2016).
- Bergstra, J., Yamins, D. & Cox, D. Making a science of model search: Hyperparameter optimization in hundreds of dimensions for vision architectures. In *International conference on machine learning*, 115–123 (PMLR, 2013).
- He, K., Zhang, X., Ren, S. & Sun, J. Identity mappings in deep residual networks. In *ECCV*, 630–645 (Springer, 2016).
- Zhang, R. Making convolutional networks shift-invariant again. In *International conference on machine learning*, 7324–7334 (PMLR, 2019).
- Li, L. *et al.* A system for massively parallel hyperparameter tuning. *Proceedings of Machine Learning and Systems 2 (MLSys20)*.
- Hochreiter, S. & Schmidhuber, J. Long short-term memory. *Neural computation* 9, 1735–1780 (1997).
- Twomey, N. *et al.* A comprehensive study of activity recognition using accelerometers. In *Informatics*, vol. 5, 27 (Multidisciplinary Digital Publishing Institute, 2018).
- Yule, G. U. On the methods of measuring association between two attributes. *Journal of the Royal Statistical Society* 75, 579–652 (1912).
- Cramér, H. *Mathematical Methods of Statistics (PMS-9)*, Volume 9 (Princeton university press, 2016).
- Efron, B. *The jackknife, the bootstrap and other resampling plans* (SIAM, 1982).
- Loshchilov, I. & Hutter, F. Sgdr: Stochastic gradient descent with warm restarts. *ICLR* (2017).
- Smith, L. N. Cyclical learning rates for training neural networks. In *2017 IEEE winter conference on applications of computer vision (WACV)*, 464–472 (IEEE, 2017).

36. Um, T. T. *et al.* Data augmentation of wearable sensor data for parkinson's disease monitoring using convolutional neural networks. In *Proceedings of the 19th ACM International Conference on Multimodal Interaction*, 216–220 (2017).
37. Chan Chang, S. *et al.* Capture-24: Activity tracker dataset for human activity recognition, <https://doi.org/10.5287/bodleian:NGx0JOMP5> (2021).
38. Doherty, A. *et al.* Gwas identifies 14 loci for device-measured physical activity and sleep duration. *Nature communications* **9**, 1–8 (2018).
39. Walmsley, R. *et al.* Reallocation of time between device-measured movement behaviours and risk of incident cardiovascular disease. *British journal of sports medicine* (2021).
40. Ordóñez, F. J. & Roggen, D. Deep convolutional and lstm recurrent neural networks for multimodal wearable activity recognition. *Sensors* **16**, 115 (2016).
41. Yuan, H. *et al.* Self-supervised learning of accelerometer data provides new insights for sleep and its association with mortality. *medRxiv* (2023).
42. Haresamudram, H. *et al.* Masked reconstruction based self-supervision for human activity recognition. In *Proceedings of the 2020 ACM International Symposium on Wearable Computers*, 45–49 (2020).
43. Saeed, A., Ozcelebi, T. & Lukkien, J. Multi-task self-supervised learning for human activity detection. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* **3**, 1–30 (2019).
44. Haresamudram, H., Essa, I. & Plötz, T. Assessing the state of self-supervised human activity recognition using wearables. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* **6**, 1–47 (2022).
45. Jain, Y., Tang, C. I., Min, C., Kawsar, F. & Mathur, A. Colloss: Collaborative self-supervised learning for human activity recognition. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* **6**, 1–28 (2022).
46. Yuan, H. *et al.* Self-supervised learning for human activity recognition using 700,000 person-days of wearable data. *npj Digital Medicine*, <https://doi.org/10.1038/s41746-024-01062-3> (2024)
47. Tong, C., Tailor, S. A. & Lane, N. D. Are accelerometers for activity recognition a dead-end? In *Proceedings of the 21st International Workshop on Mobile Computing Systems and Applications*, 39–44 (2020).
48. Weiss, G. M., Yoneda, K. & Hayajneh, T. Smartphone and Smartwatch-Based Biometrics Using Activities of Daily Living. *IEEE Access* **7**, 133190–133202, <https://doi.org/10.1109/ACCESS.2019.2940729> (2019).
49. Baños, O. *et al.* A benchmark dataset to evaluate sensor displacement in activity recognition. In *Proceedings of the 2012 ACM Conference on Ubiquitous Computing*, 1026–1035, <https://doi.org/10.1145/2370216.2370437> (ACM, Pittsburgh Pennsylvania, 2012).
50. Small, S. R. *et al.* Self-Supervised Machine Learning to Characterize Step Counts from Wrist-Worn Accelerometers in the UK Biobank. *MSSE*, <https://doi.org/10.1249/mss.0000000000003478> (2024).
51. Hoelzemann, A., Romero, J. L., Bock, M., Laerhoven, K. V. & Lv, Q. Hang-Time HAR: A Benchmark Dataset for Basketball Activity Recognition Using Wrist-Worn Inertial Sensors. *Sensors* **23**, <https://doi.org/10.3390/s23135879> (2023).
52. Berlin, E. & Van Laerhoven, K. Detecting leisure activities with dense motif discovery. In *Proceedings of the 2012 ACM Conference on Ubiquitous Computing*, 250–259, <https://doi.org/10.1145/2370216.2370257> (ACM, Pittsburgh Pennsylvania, 2012).
53. Mattfeld, R., Jesch, E. & Hoover, A. A new dataset for evaluating pedometer performance. In *2017 IEEE International Conference on Bioinformatics and Biomedicine (BIBM)*, 865–869, <https://doi.org/10.1109/BIBM.2017.8217769> (IEEE, Kansas City, MO, 2017).
54. Scholl, P. M., Wille, M. & Van Laerhoven, K. Wearables in the wet lab: A laboratory system for capturing and guiding experiments. In *Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing*, 589–599, <https://doi.org/10.1145/2750858.2807547> (ACM, Osaka Japan, 2015).
55. Sztyley, T. & Stuckenschmidt, H. On-body localization of wearable devices: An investigation of position-aware activity recognition. In *2016 IEEE International Conference on Pervasive Computing and Communications (PerCom)*, 1–9, <https://doi.org/10.1109/PERCOM.2016.7456521> (IEEE, Sydney, Australia, 2016).
56. Brunner, G., Melnyk, D., Sigfússon, B. & Wattenhofer, R. Swimming style recognition and lap counting using a smartwatch and deep learning. In *Proceedings of the 23rd International Symposium on Wearable Computers*, 23–31, <https://doi.org/10.1145/3341163.3347719> (ACM, London United Kingdom, 2019).
57. Bock, M., Kuehne, H., Van Laerhoven, K. & Moeller, M. WEAR: An Outdoor Sports Dataset for Wearable and Egocentric Activity Recognition, <https://doi.org/10.48550/ARXIV.2304.05088> (2023).
58. Yan, Y. *et al.* Topological Nonlinear Analysis of Dynamical Systems in Wearable Sensor-Based Human Physical Activity Inference. *IEEE Transactions on Human-Machine Systems* **53**, 792–801, <https://doi.org/10.1109/THMS.2023.3275774> (2023).
59. Ciliberto, M., Fortes Rey, V., Calatroni, A., Lukowicz, P. & Roggen, D. Opportunity++: A Multimodal Dataset for Video- and Wearable, Object and Ambient Sensors-Based Human Activity Recognition. *Frontiers in Computer Science* **3**, 792065, <https://doi.org/10.3389/fcomp.2021.792065> (2021).
60. Roggen, D. *et al.* Collecting complex activity datasets in highly rich networked sensor environments. In *2010 Seventh International Conference on Networked Sensing Systems (INSS)*, 233–240, <https://doi.org/10.1109/INSS.2010.5573462> (IEEE, Kassel, Germany, 2010).
61. Yang, A. Y., Kuryloski, P. & Bajcsy, R. WARD: A Wearable Action Recognition Database (2009).
62. Reiss, A. & Stricker, D. Introducing a New Benchmarked Dataset for Activity Monitoring. In *2012 16th International Symposium on Wearable Computers*, 108–109, <https://doi.org/10.1109/ISWC.2012.13> (IEEE, Newcastle, United Kingdom, 2012).
63. Frade, F. D. I. T. *et al.* Guide to the Carnegie Mellon University Multimodal Activity (CMU-MMAC) Database. Tech. Rep. CMU-RI-TR-08-22, Carnegie Mellon University, Pittsburgh, PA (2008).
64. Zappi, P. *et al.* Activity Recognition from On-Body Sensors: Accuracy-Power Trade-Off by Dynamic Sensor Selection. In Verdone, R. (ed.) *Wireless Sensor Networks*, vol. 4913, 17–33, [https://doi.org/10.1007/978-3-540-77690-1\\_2](https://doi.org/10.1007/978-3-540-77690-1_2) (Springer Berlin Heidelberg, Berlin, Heidelberg, 2008).
65. Bruno, B., Mastrogiovanni, F., Sgorbissa, A., Vernazza, T. & Zaccaria, R. Analysis of human behavior recognition algorithms based on acceleration data. In *2013 IEEE International Conference on Robotics and Automation*, 1602–1607, <https://doi.org/10.1109/ICRA.2013.6630784> (IEEE, Karlsruhe, Germany, 2013).
66. Banos, O. *et al.* mHealthDroid: A Novel Framework for Agile Development of Mobile Health Applications. In Pecchia, L., Chen, L. L., Nugent, C. & Bravo, J. (eds.) *Ambient Assisted Living and Daily Activities*, vol. 8868, 91–98, [https://doi.org/10.1007/978-3-319-13105-4\\_14](https://doi.org/10.1007/978-3-319-13105-4_14) (Springer International Publishing, Cham, 2014).
67. Chen, C., Jafari, R. & Kehtarnavaz, N. UTD-MHAD: A multimodal dataset for human action recognition utilizing a depth camera and a wearable inertial sensor. In *2015 IEEE International Conference on Image Processing (ICIP)*, 168–172, <https://doi.org/10.1109/ICIP.2015.7350781> (IEEE, Quebec City, QC, Canada, 2015).
68. Ofli, F., Chaudhry, R., Kurillo, G., Vidal, R. & Bajcsy, R. Berkeley MHAD: A comprehensive Multimodal Human Action Database. In *2013 IEEE Workshop on Applications of Computer Vision (WACV)*, 53–60, <https://doi.org/10.1109/WACV.2013.6474999> (IEEE, Clearwater Beach, FL, USA, 2013).
69. Altun, K., Barshan, B. & Tunçel, O. Comparative study on classifying human activities with miniature inertial and magnetic sensors. *Pattern Recognition* **43**, 3605–3620, <https://doi.org/10.1016/j.patcog.2010.04.019> (2010).
70. Chen, C., Jafari, R. & Kehtarnavaz, N. UTD Multimodal Human Action Dataset (UTD-MHAD) Kinect V2 (2015).

## Acknowledgements

This work is supported by Novo Nordisk (HY, SC, AD); the Wellcome Trust [223100/Z/21/Z] (AD); GlaxoSmithKline (AA); the British Heart Foundation Centre of Research Excellence [RE/18/3/34214] (AD); the National Institute for Health Research (NIHR) Oxford Biomedical Research Centre (AD); and Health Data Research UK (RW, AD), an initiative funded by UK Research and Innovation, Department of Health and Social Care (England) and the devolved administrations, and leading medical research charities. It is also supported by the UK's Engineering and Physical Sciences Research Council (EPSRC) with grants EP/S001530/1 (the MOA project) and EP/R018677/1 (the OPERA project) and the European Research Council (ERC) via the REDIAL project (Grant Agreement ID: 805194), and industrial funding from Samsung AI. The data collection was carried out using funding (JG) from the UK Economic and Social Research Council (grant number ES/L011662/1) and the European Research Council Advanced Grant (Grant number 339703). Finally, we would like to thank Rosemary Walmsley for her contributions in the dataset curation and analysis. For the purpose of open access, the author has applied a CC-BY public copyright licence to any author accepted manuscript version arising from this submission.

## Author contributions

S.C., H.Y., C.T., and A.D. conceived the experiments. S.C., H.Y., C.T., A.D., A.S. conducted the experiments. S.C., H.Y., C.T., A.D., A.A., and A.S. analysed the results. H.Y., S.C. wrote the first draft of the manuscript. All authors reviewed the manuscript.

## Competing interests

The authors declare no competing interests.

## Additional information

**Supplementary information** The online version contains supplementary material available at <https://doi.org/10.1038/s41597-024-03960-3>.

**Correspondence** and requests for materials should be addressed to A.D.

**Reprints and permissions information** is available at [www.nature.com/reprints](http://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 licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence 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 licence, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2024