IET Journals

The Institution of Engineering and Technology

Energy Expenditure Estimation using Visual and Inertial Sensors

Lili Tao[1,2], Tilo Burghardt[1],Majid Mirmehdi[1]*, Dima Damen[1], Ashley Cooper[1], Massimo Camplani[1], Sion Hannuna[1], Adeline Paiement[1], Ian Craddock[1]

[1] *SPHERE, Faculty of Engineering, University of Bristol, Bristol, UK*
[2] *Centre for Machine Vision, Bristol Robotics Laboratory, University of the West of England, Bristol, UK*
* E-mail: M.Mirmehdi@bristol.ac.uk

Abstract: Deriving a person's energy expenditure accurately forms the foundation for tracking physical activity levels across many health and lifestyle monitoring tasks. In this work, we present a method for estimating calorific expenditure from combined visual and accelerometer sensors by way of an RGB-Depth camera and a wearable inertial sensor. The proposed individual-independent framework fuses information from both modalities which leads to improved estimates beyond the accuracy of single modality and manual metabolic lookup table (MET) based methods. For evaluation, we introduce a new dataset called *SPHERE_RGBD+Inertial_calorie*, for which visual and inertial data is simultaneously obtained with indirect calorimetry ground truth measurements based on gas exchange. Experiments show that the fusion of visual and inertial data reduces the estimation error by 8% and 18% compared to the use of visual only and inertial sensor only, respectively, and by 33% compared to a MET-based approach. We conclude from our results that the proposed approach is suitable for home monitoring in a controlled environment.

1 Introduction

The term "energy expenditure" refers to a human's calorific uptake over time, which is one commonly used single metric to quantify physical activity levels. It is an important determinant in understanding the development of chronic diseases, such as obesity and diabetes. Current evidence-based guidelines [1] indicate that people who are regularly physically active have a 20% to 40% lower risk of developing conditions such as cardiovascular disease and type 2 diabetes than those who are inactive, and suggest that adults should accumulate at least 150 minutes of moderate intensity physical activity each week or 75 minutes of vigorous activity, or a combination of the two. Most research into estimating and understanding calorific expenditure focuses on coarse energy totals over longer time segments or relates to specific activities only, such as walking and running, which generally occur outside the home.

Yet, very little attention has been paid on how activities of normal daily living in an indoor environment can be quantified and understood in terms of energy expenditure. Traditionally, physical activity levels have been measured in metabolic equivalents of task (MET) [2], where a fixed value is assigned to each activity, e.g. 1 MET corresponds to energy expended at rest. However, the method is highly unreliable due to the fact that the activities are monitored using self-report approaches, such as questionnaires and occasional clinical check-ups.

There are various approaches that reliably estimate human energy expenditure via respiratory gas analysis, including both direct and indirect methods. Direct calorimetry measures, such as a sealed respiratory chamber [3], produce accurate outputs, but require a laboratory-based environment. Indirect calorimetry, on the other hand, measures energy expenditure based on inspired and expired respiratory gas flows, volumes and concentrations of oxygen consumption and carbon dioxide production. Some of these measurement devices are portable, less invasive and can produce accurate readings. They form the measurement standard for non-stationary scenarios where the person can move freely. Nevertheless, participants in experiments are required to carry gas sensors and wear a breathing mask [4].

Recently, with an increasing number of wearable devices becoming available, approximating the energy expenditure using inertial sensors has become a popular monitoring choice due to its low cost, low energy consumption, and data simplicity. Acceleration reflects a relation between motion and energy expenditure, thus tri-axial accelerometers are the most broadly used inertial sensors [5]. Recent studies show that more sensors could be involved in the task: breath rate, chest and arm skin temperature also show the correlation with energy expenditure via estimating the oxygen consumption [6]. The data could be obtained by a heart rate monitor and thermometers.

Vision-based systems, as alternative approximative sensors, do not require the wearing of extra devices. In fact, they are already a key part of home entertainment systems today [7], where RGB-Depth sensors allow for a rich and fine-grained analysis of human activity for purposes such as gaming within the field of view. Recent advances in computer vision have now opened up the possibility of integrating these devices seamlessly into home monitoring and assisted living systems [8–10].

Simultaneous utilisations of visual and inertial sensors are not common today, but are receiving growing attention in various areas, including action recognition [5], gesture recognition [11], robotics [12], augmented/virtual reality [13], and assistive technologies applications, such as fall detection [14], food preparation [15] and in a general ambient assisted living system [16]. Although employing multi-modal sensors has the advantage of complementing shortcomings of individual modalities, wearing a multitude of sensors can cause user acceptance issues.

With this in mind, in this paper we propose a framework for estimating energy expenditure in living environments based on a non-intrusive RGB-Depth visual sensor and two inertial sensors - worn on the wrist and waist - backed up in experiments by simultaneously taken indirect calorimetry measurements based on the measurement of oxygen consumption and carbon dioxide production for an accurate ground truth provision. This is a new application and to the best of our knowledge no dataset of a similar setup with reliable and accurate ground truth exists. Thus, in order to quantify the performance of the proposed system, we present a new dataset, *SPHERE_RGBD+Inertial_calorie*, for calorific expenditure estimation collected within a home environment. The dataset contains 11 common household activities performed over up to 20 sessions, lasting up to 30 minutes for each session, in each of which the activities are performed continuously. The experimental setup consists of an

RGB-Depth Asus Xtion camera mounted at the corner of a living room, two accelerometer sensors and a COSMED K4b2 [4] indirect calorimeter for ground truth measurement (see Figure 3). The *SPHERE_RGBD+Inertial_calorie* dataset is publicly released *.

This article is built upon our recent work in [17–19], with significant extensions and improvements. [17] introduced a fusion framework for recognising human daily activity using visual and inertial sensors. The work did not address the issue of energy expenditure estimation. [19] proposed a framework for calorific expenditure estimation using only a visual sensor, thus, there was no sensor fusion involved. In [18], we presented a system which allows real-time prediction for activity intensity levels, relying on light-weight bounding box features. This makes the method unable to produce precise calorific expenditure values. In this work, we have improved the feature representation for both inertial and visual sensor data by considering spatial and temporal information at the same time, and investigated both early and late fusion approaches of the data from these sensors. The key contributions of this work are: *(a)* We propose a first-ever framework for the estimation of calorific expenditure from a RGB-Depth sensor and inertial wearable sensors. There is no work published on visual-inertial energy expenditure estimation, there being only very few works that offer purely vision-based estimation [18, 20, 21]. *(b)* We improve the feature representation for both inertial and visual features in the previous fusion framework in [17] by extracting rich, multi-level information to give improved estimation accuracy. *(c)* We introduce a new dataset, linking more than 10 hours of RGB-Depth video data and inertial sensor data to ground truth calorie readings from indirect calorimetry based on gas exchange. *(d)* We present a comparative study on the utility of both visual and inertial data when estimating energy expenditure in a living environment. The visual sensor and inertial sensors are evaluated individually first, followed by an evaluation of two fusion approaches.

The remainder of the paper is organized as follows: Section 2 presents the background and work related to our study. Section 3 describes the proposed framework for estimating energy expenditure from RGB-Depth and inertial sensors alone, as well as in fusion. The experimental setup and the results are presented in Section 4, followed by a discussion and our conclusions in Section 5.

2 Related Work

2.1 Inertial Sensors

Acceleration, angular velocity, and rotation signals from wearable devices have been used for human action recognition [22], and are popular in healthcare-oriented applications, such as in fall detection systems [23] and medication adherence monitoring systems [24]. Inertial sensors can offer particularly low-cost and ubiquitous monitoring solutions for physical activities. Techniques that can control computational complexity, power consumption, and improve the unobtrusiveness of the wearable computers [25] are applicable to many systems including the one at hand. Here, we firstly discuss inertial sensor feature extraction methods described in the literature, followed by an outline of existing models of energy expenditure estimation based on them.

Feature Representation. Different features extracted from inertial sensor devices have been considered ranging from raw signal samples to high-level descriptors. Raw time series data from accelerometers is most often provided as triples of scalars, where each scalar corresponds to acceleration in one of three orthogonal spatial dimensions. The same fundamental structure applies to angular velocity signals and orientation signals of three directions. There is no computational burden associated with feature extraction when the raw data is used.

However, raw data may not expose enough discriminative structure to achieve high performance on specific classification tasks.

Instead, statistical features may be extracted from each of the three axes, where sensor signal sequences are often partitioned into temporal windows over which features are generated. All features extracted from a temporal window are then concatenated to form a single combined descriptor vector.

Commonly used features include the first- and second-order statistics, namely the mean and variance [26]. In [17], apart from these commonly used features, correlation measures between each axes pair are also extracted. Basic statistical measures are computationally efficient and are able to capture structural patterns in inertial data. The feature descriptor can be further quantised into a number of codewords, such as in [27]. Approaches based on deep learning are currently being explored to create more generalised learning methods that generate features directly from the input data and promise to optimise performance further [28].

Energy Expenditure Estimation. The first automatic methods for inertial energy expenditure estimation [29] were count-based estimation systems applied by fitting a single regression model to all the data regardless of what activity was being performed. However, systems that map from a single wearable to calorie values struggle to accurately estimate the intensity of physical activity across a range of actions. For example, some actions involving only upper or lower body movements are difficult to be recognised via a single wearable device, and therefore a high estimation error [30] occurs. Different activities may require different models to represent them. Activity-specific methods split the estimation process into two steps, where activity groups are detected and classified first, and only then an activity-specific model is applied to estimate energy expenditure. MET lookup tables are the most common approach to perform the latter, where a static MET value is assigned from a compendium on physical activities [2] to each one of the clusters of activities [31]. However, METs-based approaches neglect any transitional effects of activities (continued calorie expenditure after rigorous activity has finished), and they overlook the fact that even the activities in same cluster can be performed at varying intensities, for example, walking at different speeds, or body exercise with different intensity.

An attempt to model the transition between activities was proposed in [32], where an accelerometer and a heart rate sensor were used and the transition between sedentary, household activities and walking were modelled. The work in [33] shows that by using data from multiple inertial sensors one can more accurately predict energy expenditure, although the limitations of wearable devices are considerable particularly with respect to accuracy as emphasised in [17].

Accelerometer feature descriptors are often formed within a temporal window. This brings out another concern that the window sizes are set usually at less than 10 seconds in existing works [32][6]. The length of window would significantly affects the results. It should be short enough to recognise activities as local temporal information are more descriptive, but long enough to predict calorie values since current energy expenditure strongly depends on previous activity intensity level.

2.2 Visual Sensors

Visual sensor based techniques have emerged over recent years for which there exists a significant body of literature describing the inference of activities from 2D colour intensity imagery [34] Meanwhile, the increasing availability of depth-measuring sensors, especially the introduction of the Microsoft Kinect, has generated an opportunity for utilizing depth in conjunction with traditional RGB camera data allowing for richer and more fine-grained analysis of human activity [7]. Applying computer vision techniques to help with the diagnosis and management of health and wellbeing conditions has gained significant momentum over the last years [35]. However, studies on energy expenditure using visual sensors have been relatively limited. Our work explores this field further and builds on several relevant subject areas in computer vision.

Visual Feature Representation. The visual trace of human activity in video forms a spatio-temporal pattern. To extract relevant properties from this for the task at hand, one aims at compactly capturing this pattern and highlighting important aspects related to the properties of interest. Assuming that both body configuration and body motion [36] are relevant to infer calorific uptake, the pool of potential features is large - ranging from local interest point configurations [37], over holistic approaches like histograms of oriented gradients and histograms of motion information [17], to convolutional neural network features [38].

Motion information in the first place could also be recovered in various ways, e.g. from RGB data using optical flow *or* from depth data using 4D surface normals [39]. Whilst a composition of these features via concatenation of per-frame descriptors is straight forward, this approach suffers from the curse of dimensionality and unaffordable computational cost. Sliding window methods [40], on the other hand, can limit this by predicting current values only from nearby data within a temporal window. Further compaction may be achieved by converting large feature arrays into a single, smaller vector with a more tractable dimension count via, for instance, bags of visual words [41], Fisher vectors [42], time series pooling [43], or the features extracted from convolutional neural networks [44]. In summary, the challenge of feature representation will require capturing visual aspects relevant to calorific expenditure, whilst limiting the dimensionality of the descriptor.

Activity Recognition. There exists a significant body of literature describing the inference of activities from 2D colour intensity imagery [34], RGB-Depth data [7], and skeleton-based data [45]. Knowledge about the type of activity undertaken has been shown to correlate with the calorific expenditure incurred [2]. In alignment with Figure 1(a), we will argue in this work that an explicit activity recognition step in the vision pipeline can, as an intermediate component, aid the visual estimate of energy uptake.

Energy Expenditure Estimation. 2D video has recently been used by Edgcomb and Vahid [20] coarsely to estimate daily energy expenditure. In their work, subjects are first segmented from the scene background. Changes in height and width of the subject's motion bounding box, together with vertical and horizontal velocities and accelerations, are then used to estimate calorific uptake. Tsou and Wu [21] take this idea further and estimate calorie consumption using full 3D joint movements tracked as skeleton models by a Microsoft Kinect. We note, however, that both of the above methods use wearable accelerometry as the target ground truth, which in fact does not provide an accurate benchmark; skeleton data is commonly noisy and currently only operates reliably when the subject is facing the camera. This limits applicability in more complex, in-the-wild, visual settings as, for instance, contained in the *SPHERE_RGBD+Inertial_calorie* dataset. Our recent work in [19] introduced a visual based framework for estimating calorific expenditure in a home environment, and we then extended it to be able to estimate physical activity intensity levels in real-time [18]. Although the method is practically applicable to more complex settings, the light-weight features extracted from bounding boxes (velocity vector and the ratio of height and width of the bounding box) can only help make a gross estimate of calorific expenditure. In this paper, instead of using only simple bounding box features, we simultaneously collect RGB and depth imagery and then encode appearance and motion features via spatial pyramids. The temporal information is encoded using a pyramidal temporal pooling with multiple pooling operators. This has the aim of extracting rich, multi-level information to give improved estimation accuracy, whilst maintaining applicability to detect complex human activities.

2.3 Sensor Fusion

It is reasonable to expect that the use of multiple sensor types would improve the overall performance compared to single sensor settings, since sensors may complement the limitations of each other. Given an accurate temporal synchronisation between the different modality sensors, learning from multi-modal data is applicable. In general,

feature-level fusion (early fusion) and decision-level fusion (late fusion) are the two approaches most often employed to fuse multiple modalities. Both early and late fusion strategies are explained in further detail in [46].

Feature-level Fusion. This methodology involves carrying out fusion of features right after features are extracted from raw data. This scheme only requires one learning stage and allows to take advantage of mutual information from data. For instance, in [47], the depth and inertial sensor data were concatenated, then an HMM classifier was employed for recognising basic gestures on the fused data. The results reveal significant improvements when the fusion scheme is applied compared to using each sensor individually. The work in [17] investigates the practical home-use of body-worn mobile phone inertial sensors together with an RGB-Depth camera to achieve monitoring of daily living scenarios. The results indicate that the vision-based approach significantly outperforms the wearable-based method, while fusion of both sensors slightly improves the performance further. Clearly, feature-level fusion can be applied effectively in practical settings, however, it may suffers from the 'curse of dimensionality'.

Decision-level Fusion. This approach fuses the decisions made by individual classifiers, each of which corresponding to one sensor. Since decision information is of low complexity, the curse of dimensionality can effectively be targeted. In [48], for instance, a Bayesian co-boosting training framework combines multiple hidden Markov model classifiers of two modalities - a Kinect sensor and an inertial measurement unit. The result is the construction of a strong classifier for gesture recognition, which achieved the best performance in the Multi-modal gesture recognition challenge. A real-time action recognition system in [49] uses Dempster-Shafer theory to combine the classification outcomes from a depth camera and several inertial sensors. A Bayesian model for sensor fusion is introduced in [16], which aims at addressing the challenges of fusion of heterogeneous sensor modalities in ambient assisted living.

Comparisons. In this work, we consider both fusion approaches and provide a direct comparison. In the feature-level fusion approach, features generated from the two modality sensors are merged before classification, and the decision-level fusion is performed by forming a linear combination of different classifiers using stacking regression [50] to improve overall accuracy. As outlined in the following section, our work attempts to use skeleton-independent, RGB-Depth-based vision, together with two wearable accelerometer devices to estimate calorific expenditure against a standardised calorimetry sensor COSMED-K4b2 based on gas exchange.

3 Method

To describe our framework for estimating calorific expenditure, we initially introduce the methods for visual and wearable sensors separately, and then describe two approaches for their fusion.

Figure 1(a) shows a flowchart of the visual method – mapping visual flow and depth features to calorie estimates using activity-specific models. The method implements a cascaded and recurrent approach, which explicitly detects activities as an intermediate to select type-specific mapping functions for the final calorific estimation. Importantly, our setup as a video-based system is designed to reason about activities *first*, before estimating calorie expenditure via a set of models which are each separately trained for particular activities. In contrast to this, our direct mapping method designed for wearable sensor data directly maps inertial features to calorie estimates via a monolithic classifier. A flowchart of the wearable approach is shown in Figure 1(b). In our fusion system we consider both feature-level and decision-level fusion of these two approaches. Finally, we compare these methods against a ground truth of gas-exchange measurements and off-the-shelf alternatives, that is manual mapping from activity classes to calorie estimates via metabolic equivalent task (METs) lookup tables [2] as it is often applied in clinical practice today.

(a) Visual-based framework

(b) Wearable-based framework
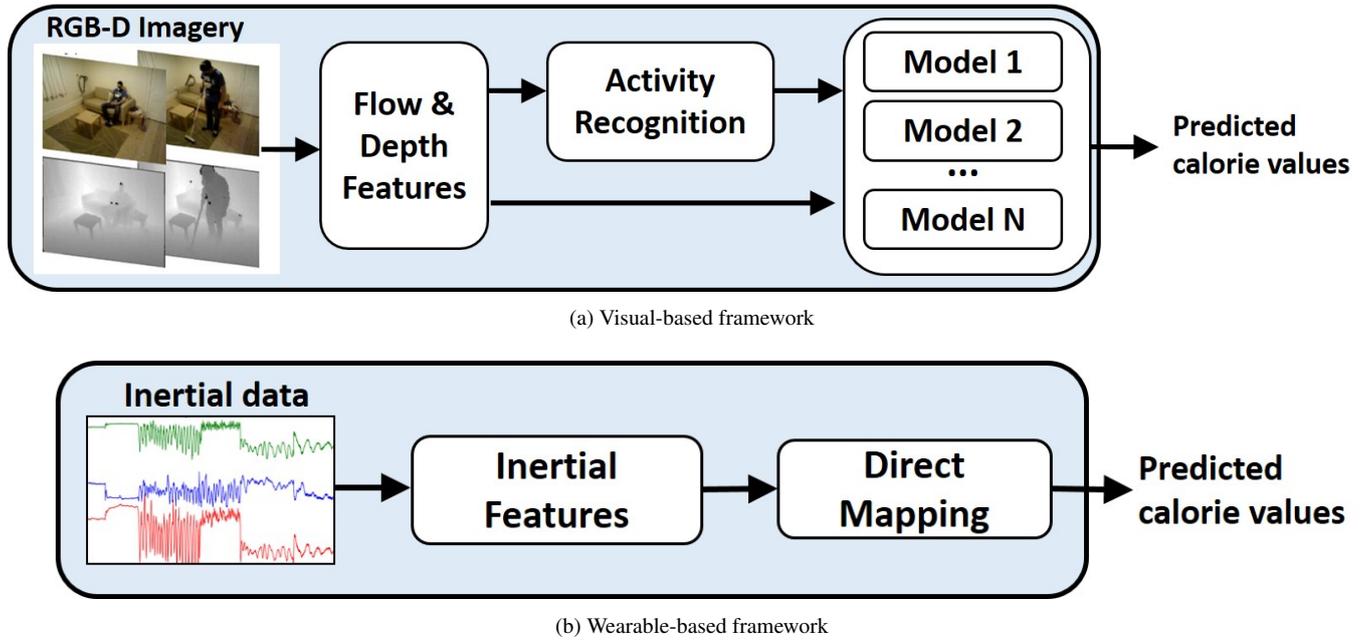
Fig. 1: **Overview of our Visual-based and wearable-based frameworks.** (a) Visual-based framework. RGB-Depth videos are represented by a combination of flow and depth features. The proposed recurrent method then selects activity-specific models which map to energy expenditure estimates. (b) Wearable-based framework. Inertial features are formed from two accelerometer sensor data, then features are mapped directly to calorie estimates via a monolithic classifier.

3.1 Visual Features

We obtain RGB and depth imagery using an Asus Xtion. For each frame t, appearance and motion features are extracted, with the latter being computed with respect to the previous frame (level 0). A set of temporal filters is then applied to form higher level motion features (level 1). We extract features from within the bounding box returned by the OpenNI SDK [51] person detector and tracker, which allows to follow up to 6 persons in the camera view at the same time. To normalise the utilised image region due to varying heights of the subjects and their distance to the camera, the bounding box is scaled by fixing its longer side to $M = 60$ pixels, a size recognised as optimal for human action recognition [52], while maintaining aspect ratio. The scaled bounding box is then centred in a $M \times M$ square box and horizontally padded.

Motion Feature Encoding - Inspired by [52], optical flow measurements are taken over the bounding box area and split into horizontal and vertical components. These are re-sampled to fit the normalised box and a median filter with kernel size 5×5 is applied to smooth the data. A spatial pyramid structure is used to form hierarchical features from this. Such partitioning of the image into an iteratively growing number of sub-regions increases discriminative power. The normalised bounding box is divided into a $n_g \times n_g$ non-overlapping grid, where n_g depends on the pyramid level, and the orientations of each grid cell are quantised into n_b bins. The parameters for our experiments are empirically determined as $n_b = 9$ and $n_g = 1$ and 2 for levels 0 and 1 respectively. Figure 2(a) exemplifies optical flow patterns and their encoding in two different example activities.

Appearance Feature Encoding - We extract depth features by applying the histogram of oriented gradients (HOG) feature on raw depth images [53] within the normalised bounding box. We then apply Principal Component Analysis and keep the first 150 dimensions of this high-dimensional descriptor, which retains 95% of the total variance.

Pyramidal Temporal Pooling - Given the motion and appearance features extracted from each frame in a sequence of images, it is important to capture both short and long term temporal changes, and summarise them to represent the motion in the video. Pooled motion features were first presented in [43], even though designed for egocentric video analysis. We modify their pooling operator to make it

more suitable for our data as follows - an illustration of the temporal pyramid structure and the process for pooling operations are shown in Figure 2(b). The time series data $\mathbf{S}$ can be represented as a set of time segments at level i as $\mathbf{S} = [\mathbf{S}_i^1, \ldots, \mathbf{S}_i^{2^i}]$. The final feature representation is a concatenation of multiple pooling operators applied to each time segment at each level. The time series data can also be explained as T per-frame feature vectors, such that $\mathbf{S} = \{S_1, \ldots, S_N\}$, $\mathbf{S} \in \mathbb{R}^{N \times T}$ for a video in matrix form, where N is the length of the per-frame feature vector, and T is the number of frames. A time series $S_n = [s_n(1), \ldots, s_n(T)]$ is the n^{th} feature across $1, \ldots, T$ frames, where $s_n(t)$ denotes n^{th} feature at frame t. A set of temporal filters with multiple pooling operators is applied to each time segment $[t_{min}, t_{max}]$ and produces a single feature vector for each segment via concatenation. We use two conventional pooling operators, max pooling and sum pooling, as well as frequency domain pooling. They are defined as:

$$\mathcal{O}_{\max}(S_n) = \max_{t = t_{\min} \cdots t_{\max}} s_n(t) \quad \text{and} \quad \mathcal{O}_{\text{sum}}(S_n) = \sum_{t = t_{\min}}^{t_{max}} s_n(t)$$
(1)

Frequency domain pooling is used to represent the time series S_n in the frequency domain by the discrete cosine transform (dct), where the pooling operator takes the absolute value of the j lowest frequency components of the frequency coefficients D,

$$\mathcal{O}_{\text{dct}}(S_n) = \left| M_{1:j} S_n \right|$$
(2)

where M is the discrete cosine transformation matrix.

3.2 Inertial Features

Raw time series data from accelerometers is measured as $[X, Y, Z]$ vectors, where each column corresponds to acceleration in orthogonal spatial dimensions. Figure 3 illustrates the raw accelerometer data collected from one wearable device for various actions. From the raw data the pooled motion features are formed from each of the three axes for each device. Abstracting short-term and long-term changes in the inertial feature descriptor is essential; it is particularly useful for modelling the level of activity intensity changes. Thus,
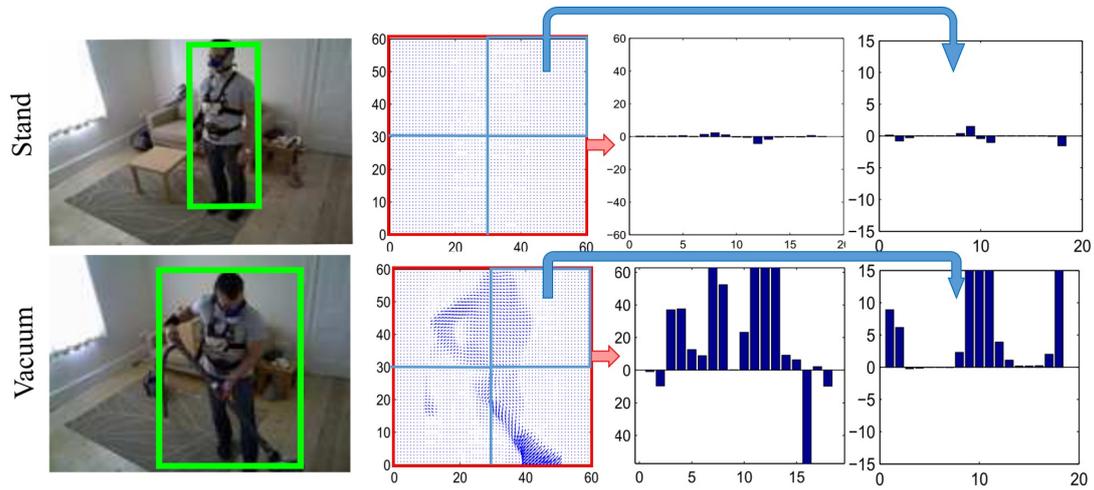
(a) Flow feature encoding via spatial pyramids
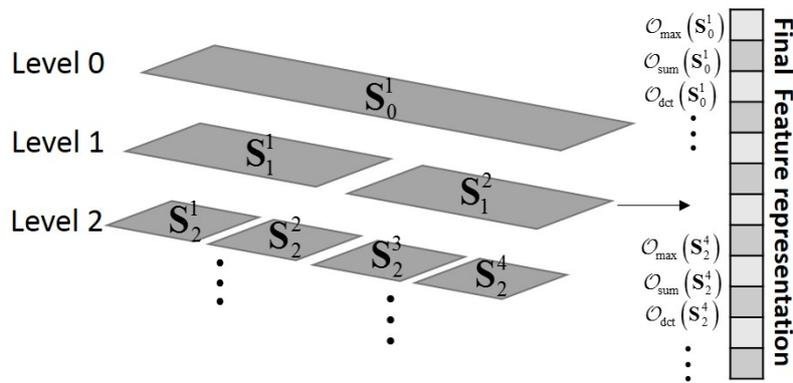
(b) Temporal pyramid pooling

Fig. 2: (a) **Flow feature encoding via spatial pyramids.** *First row:* limited motion while standing still. *Second row:* significant motion features when moving during vacuuming. *First column:* colour images with detected person. *Second column:* optical flow patterns. *Third column:* motion features at level 0. *Last column:* motion features from the top-right quadrants of the image at level 1 (at which the image is subdivided into four quadrants). (b) **Temporal pyramid pooling and its feature representation.** This schematic shows the temporal subdivision of data into various pyramidal levels (left) and the concatenation of resulting feature (e.g. max, sum and dct) into a descriptor vector (right).

we apply three pooling operators (max pooling, sum pooling and frequency domain pooling) to the inertial data.

3.3 Learning and Recurrency

Energy expenditure estimation can be formulated as a sequential and supervised regression problem. We train a support vector regressor to predict calorie values from given features over a training set. The sliding window method is used to map each input window of width w to an individual output value y_t. The window contains the current and the previous $w - 1$ observations. The window feature is represented by temporal pooling from the time series $\mathbf{S} = \left\{ S^{t-w+1}, \ldots, S^t \right\}$.

We note that energy values for a particular time are highly dependent on the energy expenditure history. In our system, these are most directly expressed by previous calorific predictions during operation. Thus, employing recurrent sliding windows offers an option to not only use the features within a window, but also take the most recent d predictions $\left\{ \hat{y}^{t-d}, \ldots, \hat{y}^{t-1} \right\}$ into consideration to help predict y^t. During learning, as suggested in [54], the ground truth labels in the training set are used in place of recurrent values.

3.4 Fusion Approach

Both feature-level and decision-level fusion are considered in our work.

Feature-level Fusion. This is an early fusion approach, for which all features from all modalities are concatenated together, and employed as a single unified feature stream to the learning components. Given visual features in d_1-dimensional feature space $\mathbf{S}_v \in \mathbb{R}^{d_1}$ and accelerometer features in d_2-dimensional feature space $\mathbf{S}_a \in \mathbb{R}^{d_2}$, the fused feature set can be represented as $\mathbf{S} \in \mathbb{R}^{d_1 + d_2}$, where the feature set is constructed as $\mathbf{S} = (\mathbf{S}_v, \mathbf{S}_a)$. The fused feature vector is then used as input to the classifiers of the system. Figure 4(a) shows a flowchart of this feature-level fusion approach.

Decision-level Fusion. In this approach, a collection of models are learned, and the predictions are combined together only at the last stage to form the final decision. We apply the decision-level fusion via a stacking regression method, which forms linear combinations of different classifiers to improve overall estimation accuracy.

Consider that there are K predicted values $\hat{y}_1, \ldots, \hat{y}_K$ estimated from each regressor individually. Then, the final predictor value $\hat{Y}(\mathbf{S})$ can be represented as a linear combination of a set of predicted values with different weighting coefficients, constructed as:

$$\hat{Y}(\mathbf{S}) = \sum_{k=1}^{K} \alpha_k \hat{y}_k(\mathbf{S}) \tag{3}$$
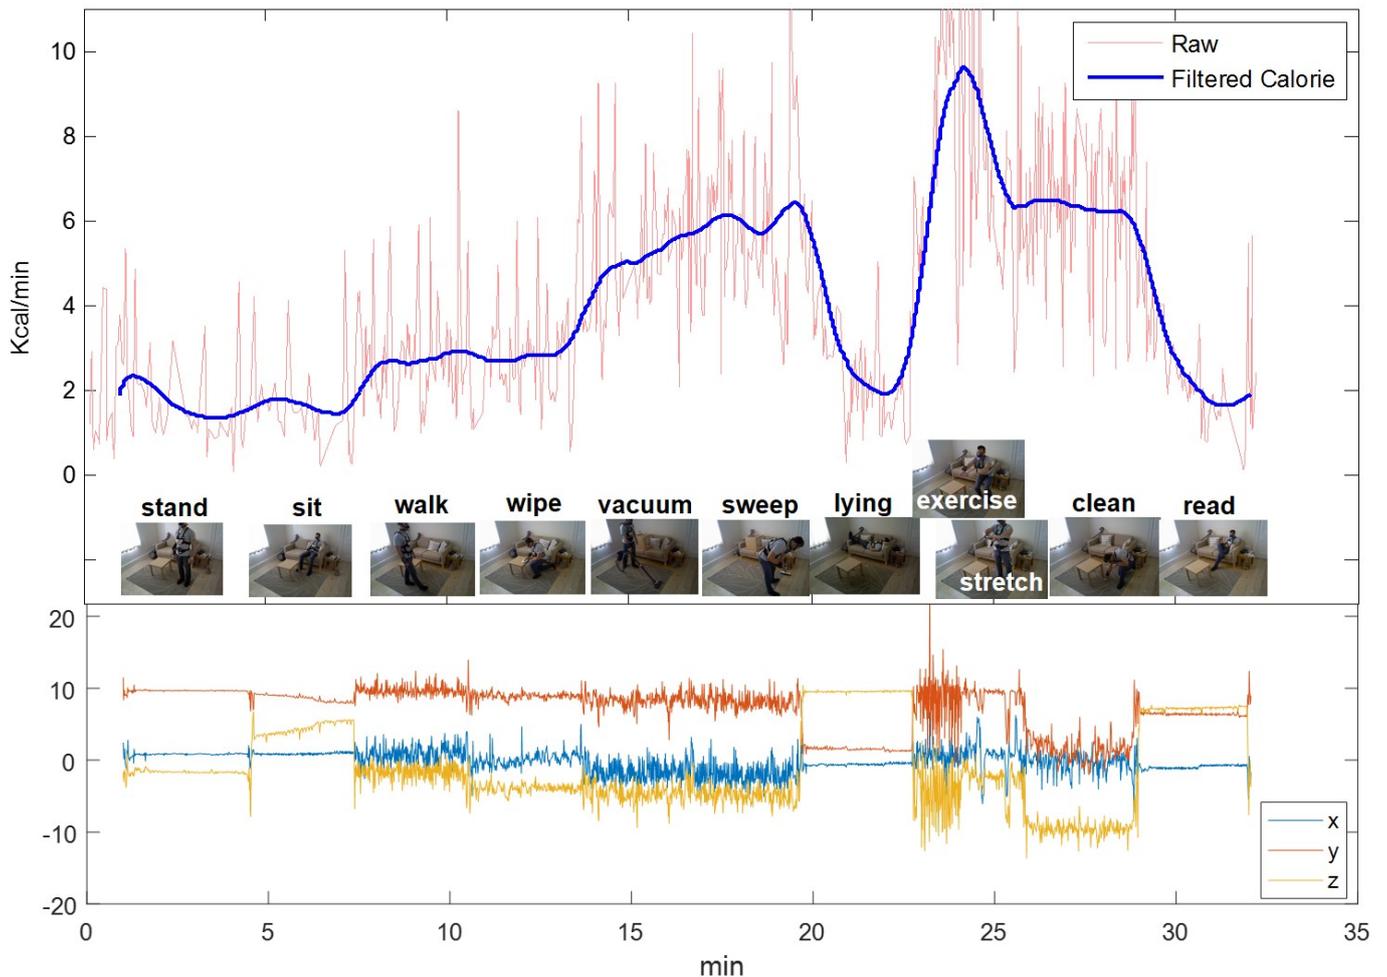
Fig. 3: **Ground truth example sequence.** Top: Raw per breath data (red) and smoothed COSMED-K4b2 calorimeter readings (blue) and sample colour images corresponding to the activities performed by the subject. Bottom: Three-axis acceleration signals from the waist-wear sensor.

Given a set of training data $\{(\mathbf{S}^1, y^1), \ldots, (\mathbf{S}^T, y^T)\}$ with T training samples, where each $\mathbf{S}^t$ is an input vector, the goal is to minimise the distance of the ground truth y^t and the predicted values $\hat{Y}^t(\mathbf{S})$ from the combined regressor. This optimised distance can be obtained by

$$\arg\min_{\alpha_k} \sum_{t=1}^{T} \left(y^t - \sum_{k=1}^{K} \alpha_k \hat{y}_k^t(\mathbf{S}^t) \right)^2, \qquad (4)$$

with the constraints $0 \leq \alpha_k \leq 1, k = 1, \ldots, K$. The resulting combined predicted value $\sum_{k=1}^{K} \alpha_k \hat{y}_k(\mathbf{S})$ is then used as prediction. Figure 4(b) shows a flowchart of this decision-level fusion approach.

## 4 Experimental Results

4.1 Dataset and Ground Truth

We introduce the *SPHERE_RGBD+Inertial_calorie* dataset for human calorific expenditure estimation, comprising RGB-Depth and inertial sensor data captured in a real living environment. The ground truth was captured by the COSMED K4b2 portable metabolic measurement system. The dataset was generated over 20 sessions by 10 subjects with varying anthropometric measurements. Participants were 7 males and 3 females, with mean age of 27.2 ± 3.8 years, mean weight of 72.3 ± 15.0 kg, mean height of 173.6 ± 9.8, mean BMI of 23.7 ± 2.8. Ethics approval was obtained and each participant signed a consent form agreeing to share their data for research

purposes. The dataset contains up to 11 activity categories per session, and totalling around 10 hours recording time. The activities were captured in daily-living scenarios containing a variety of body positions, view-points and distances naturally associated with the various actions performed. Figure 5 shows frames from the *vaccuming* activity depicting this variety. It is also shown that the sequences are captured in different time of the day which contains various lighting conditions. All the activities, the associated intensity categories and MET values are shown in Table 1. In addition, Table 2 lists the number of frames for each action and sequence[*].

Colour and depth images were acquired at a rate of 30Hz. The accelerometer data was captured at about 100Hz and sampled down to 30Hz, a frequency recognised as optimal for human action recognition [55]. The calorimeter gives readings per breath, which occurs approximately every 3 seconds. To model transitions better between activity levels, we consider the 9 different combinations of the three activity intensities (Light, Light+, Moderate) in the design of each session.

Figure 3 shows a detailed example of calorimeter readings and associated sample RGB images from the dataset (top) and the accelerometer data reading (bottom). The raw breath data is noisy (in red). We apply an average filter with a span of approximately 20 breaths (in blue). The participants were asked to perform the

---

[*]*Some actions in certain sequences are missing due to various reasons (hence they have 0 frames), for example exercise is missing in sequences 7 and 17 as the participants had difficulty in performing the action.*
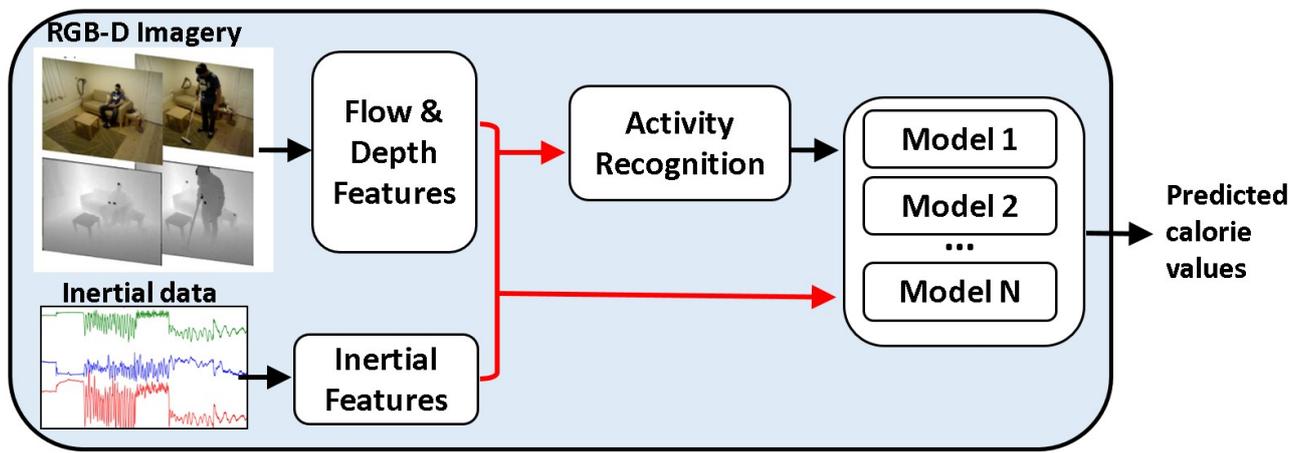
(a) Feature-level fusion
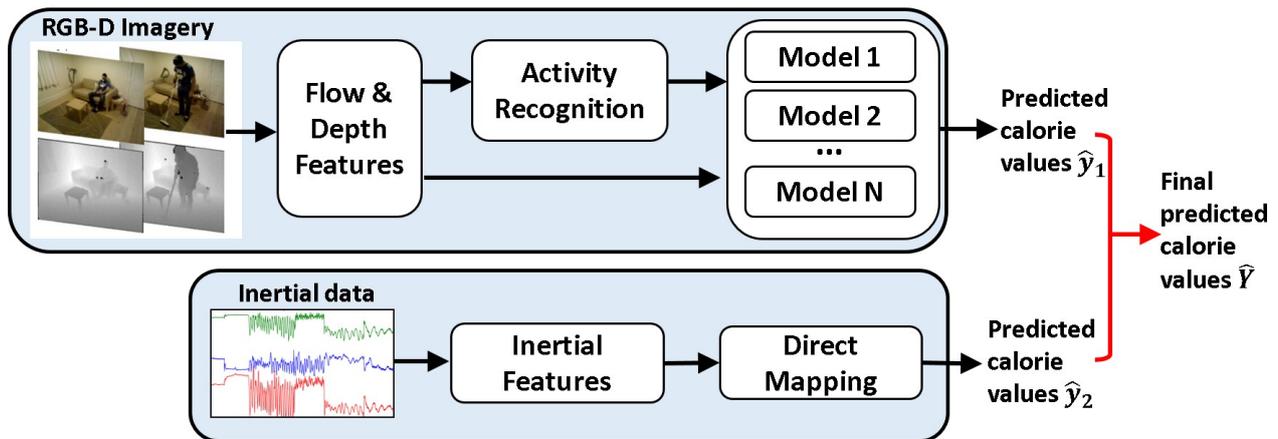
(b) Decision-level fusion

Fig. 4: **Fusion approaches overview.** (a) Feature-level fusion framework. The features from visual and inertial sensors are concatenated to form a monolithic input into activity recognition and activity-specific models. (b) Decision-level fusion framework. Calorie values are predicted individually by the different sensor modalities, and then combined using a regression method to form final calorie estimates
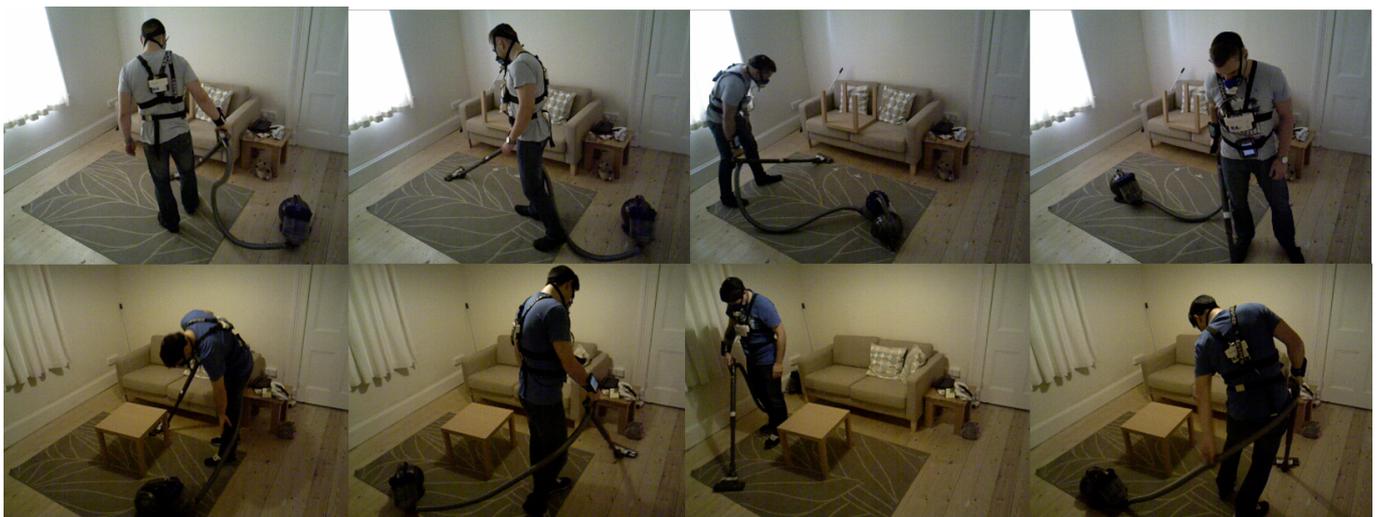
Fig. 5: Example poses from the activity "vacuuming". It can be seen that the sequences contain a large variety of body positions, view-points and distances naturally associated with the action. Two example sequences are captured in daytime and nighttime respectively indicating different lighting conditions.

activities based on their own living habits without any extra instructions.

seq ID	stand	sit	walk	wipe	vacuum	sweep	lying	exercise	stretch	clean	read	overall
1	3739	5230	5509	5117	5447	5113	5142	2789	2421	5072	6092	51671
2	3546	5388	5304	5244	5177	5091	5020	2545	2566	5112	5311	50304
3	3876	5682	6093	4745	5151	4946	5360	2297	2318	5317	5095	50880
4	3948	5211	5470	5182	5025	4813	3784	0	0	0	0	33433
5	3239	5133	5670	4878	5436	4327	0	0	0	0	0	28683
6	3812	5294	5889	5023	5015	4841	4878	3205	1778	4738	5070	49543
7	3796	5239	11602	6684	4432	4561	3590	0	5474	5837	5337	56552
8	3951	5257	5412	5244	4886	1005	0	0	0	0	0	25755
9	4128	5568	5091	5195	4651	3619	4891	2458	2483	5700	5298	49082
10	3649	5202	5317	5354	4651	4990	5030	2330	1669	3933	5337	47462
11	4367	4901	5503	5133	5006	4761	0	0	0	0	0	29671
12	3697	5270	5618	5010	5107	4988	4991	2891	2335	5299	5412	50618
13	4250	5936	4644	5162	5259	4517	4944	899	2839	5495	6206	50151
14	4263	4732	5370	5150	4769	4847	4861	3008	3224	4366	5765	50355
15	3457	5784	4789	4745	5159	4911	0	0	0	0	0	28845
16	3919	5466	5062	5308	2716	0	0	0	0	0	0	22471
17	3613	5343	5432	4914	5032	4461	4979	0	5063	5902	4873	49612
18	3715	5340	5422	5013	5893	4743	4517	1977	2793	5948	5012	50373
19	4521	5434	5787	4740	5015	4459	5480	3174	2707	4803	6342	52462
20	4040	5255	5597	5472	5309	4551	4926	1797	1691	6356	6443	51437

Table 2 Number of frames per sequence and action in *SPHERE_RGBD+Inertial_calorie* dataset

Intensity	Activity	MET value
	sit still	1.3
Light	*stand still*	1.3
	lying down	1.3
	reading	1.5
	walking	2.0
Light+	*wiping table*	2.3
	cleaning floor stain	3.0
	vacuuming	3.3
Moderate	*sweeping floor*	3.3
	upper body exercise	4.0
	stretch	5.0

Table 1 Activities, their associated MET values and the levels of activity intensity.

4.2 Parameter Settings

In our experiments, we use non-linear SVMs with Radial Basis Function kernels for activity classification and a linear support vector regressor for energy expenditure prediction. The libsvm [56] implementation was used. We perform a grid search algorithm to estimate the hyper-parameters of the SVM. To test our individual-independent approach, we implement leave-one-subject-out cross validation on the dataset in which each subject's data are tested in turn using models trained with all other subject data combined. This process iterates through all subjects, and the average testing error and standard deviation of all iterations are reported. We use the normalised root-mean-squared error (normalised RMSE) as a standard evaluation metric to facilitate the comparison between data with different scales for the deviation of estimated calorie values from the ground truth.

4.3 Evaluation of Individual Modalities

We start with tests on each sensor type (visual and inertial), and compare their performance in situations when used independently.

Temporal window size. The accuracy of predicted calorie values is linked to the window of previous information utilised for making the prediction. In a first experiment we look at the relation between window length on the one hand, and activity recognition and calorie prediction errors on the other. All sequences are tested with temporal windows of $w = \{7.5, 15, 30, 60\}$ seconds. Tables 3 and 4 illustrate the activity recognition rates and the average normalised RMSEs for calorie prediction at different window length w using visual and accelerometer data, respectively. It can be seen that, in both modalities, the best performance across the set for recognising activities is achieved when a relatively small size of window is applied. The confusion matrices corresponding to the use of visual and inertial sensors are depicted in Figure 6).

In a second experiment, we test how the estimated calorie value is influenced by the performance of action recognition. The results for each activity using visual and inertial sensors are listed in Tables 5 and 6, respectively. We first test a system in which the ground truth labels are used to select the activity-specific model for calorie prediction (top rows in Tables 5 and 6). We then compare actual action recognition at varying window lengths $w = \{7.5, 15, 30, 60\}$ seconds. In all cases, we use a fixed window length $w = 60$ seconds for calorific expenditure estimation to focus on the effect of varying action recognition quality. As expected, it can be observed that for most activities, the calorie estimation error is smallest when there is no activity recognition error (top rows in Tables 5 and 6). For a more detailed visualisation, we also show the results in Figure 7; the 11 actions are grouped into three clusters based on their intensity level (see Table 1). The figure summarises the calorie prediction error for different intensities and action recognition rates using the visual system and the inertial system, respectively.

Model comparison. As just observed, activity recognition accuracy affects the calorie prediction results. To determine if the activity-specific model provides a predictive advantage, we compare the estimation performance of the activity-specific (AS) approach against the direct mapping (DM) approach for each sensor modality. For both sensor systems, we select a fixed window length of $w = 60$ seconds to analyse performance for calorie value prediction in both DM and AS, and $w = 15$ seconds for activity recognition in AS. The results are shown in Figure 8(a) for each activity. It can be seen that for the visual-based system, the AS provides best prediction results overall and significantly outperforms DM in most activities. For the inertial-based system, the error associated with AS is significantly higher compared to DM. This is in part due to poor activity recognition results in an inertial measurement setup, which effectively leads to using wrong models to estimate calorie values.

Evaluation of a recurrent system layout. To evaluate the use of recurrency, we set the AS method using the sliding window technique as our baseline method for the vision-based comparison,

w		stand	sit	walk	wipe	vacuum	sweep	lying	exercise	stretch	clean	read	**overall**
7.5	activity	83.5	75.2	**90.5**	70.2	79.7	73.4	58.6	36.6	54.1	**93.8**	39.9	72.3
	calorie	0.84	0.70	**0.29**	0.52	**0.31**	0.41	0.74	0.63	0.52	0.41	0.67	0.58
15	activity	**86.5**	77.6	88.3	69.4	79.0	**76.5**	62.3	39.2	**61.1**	91.4	38.9	**73.7**
	calorie	0.83	0.66	0.30	0.46	0.34	0.45	0.74	0.66	0.54	0.40	0.64	0.53
30	activity	85.0	79.1	89.4	**71.9**	**81.1**	75.2	54.3	**40.3**	57.8	90.4	36.8	71.1
	calorie	0.73	0.52	**0.30**	**0.41**	0.36	0.41	0.66	0.55	0.54	**0.37**	0.54	0.49
60	activity	81.1	**79.7**	85.1	66.0	77.2	72.9	33.0	29.3	52.7	90.0	35.9	68.2
	calorie	**0.54**	**0.45**	0.32	0.44	**0.34**	**0.39**	**0.58**	**0.42**	**0.52**	0.38	**0.50**	**0.45**

Table 3 **Vision-based prediction results.** Activity recognition rate (%) and calorific expenditure prediction error (normalised RMSE) with different window lengths w, stated in seconds.

w		stand	sit	walk	wipe	vacuum	sweep	lying	exercise	stretch	clean	read	**overall**
7.5	activity	**91.1**	58.7	78.1	43.6	50.1	**48.2**	66.0	51.8	65.2	59.8	48.0	60.1
	calorie	0.54	0.50	0.49	0.52	0.51	0.41	0.64	0.83	0.72	0.51	0.77	0.63
15	activity	89.7	**58.9**	**86.7**	43.9	**50.5**	37.9	63.4	**66.7**	**68.1**	66.6	**61.8**	**60.3**
	calorie	0.50	0.51	**0.41**	**0.45**	**0.46**	0.38	**0.53**	0.80	0.68	0.44	0.73	0.59
30	activity	88.7	56.9	83.9	46.6	48.7	36.4	69.4	51.8	64.1	61.0	52.1	58.0
	calorie	0.46	0.49	0.48	0.46	0.52	**0.35**	0.54	0.77	0.60	0.43	0.71	0.56
60	activity	90.8	55.3	74.8	**52.3**	49.6	45.7	**73.1**	44.9	59.0	64.2	56.0	58.3
	calorie	**0.41**	**0.41**	0.51	0.50	0.47	0.40	0.57	**0.64**	**0.55**	**0.41**	0.71	**0.54**

Table 4 **Inertial-based prediction results.** Activity recognition rate (%) and calorific expenditure prediction error (normalised RMSE) with different window lengths w, stated in seconds.

and the DM method for inertial-based comparison. We now introduce two methods, which are based on recurrent sliding window approaches. The first method (Recurrent1) uses the most recent predictions of the baseline method as input together with visual/inertial features to predict current calorie value. Thus, it implements indirect recurrency utilising the predicted values from the baseline as recent predictions. The second method (Recurrent2) implements full recurrency, i.e. it uses its own output as recurrent input together with visual/inertial features.

Table 7 shows the effect of using recurrent information, with the best results for each activity highlighted. In general, the full recurrency model, Recurrent2, suffers from drift and produces the worst results for half of the activities and also overall. When the visual sensor is used, indirect recurrency, Recurrent1, outperforms the other approaches at an average normalised RMSE of 0.42, while in inertial-based systems indirect recurrency increases the estimation error by 7% comparing to its baseline.

4.4 Comparing Sensor Fusion Approaches

Having tested the two modalities individually, we now study modality fusion approaches against the use of individual sensor systems, and also compare against the MET look-up table method for completeness.

In feature level fusion, we apply our AS approach with $w = 15$ seconds for activity recognition and $w = 60$ seconds for calorific expenditure estimation. In decision level fusion, we again use the most suitable model for each sensor data to fuse, which is the AS approach for visual sensor data, and DM approach for inertial sensor data. The estimation performance of the two fusion approaches is compared with the performance of each sensor modality individually, as shown in Figure 8(b). It can be seen that both fusion approaches on average outperform unimodal prediction. In particular, by combining the features from the visual and the inertial data, the overall prediction error decreases from 0.46 (inertial sensors alone) and 0.42 (visual sensor alone) to 0.39. The calorie prediction accuracy for most activities are improved when using fusion approaches. We also observed that the two fusion frameworks achieve similar performance.

Finally, we present the results produced by MET, which is commonly used by clinicians and physiotherapists, and compare our

proposed methods against it. It assumes N clusters of activity $\mathbf{A} = \{A_1, A_2, \ldots, A_N\}$ are known. A MET value is assigned to each cluster, together with anthropometric characteristics of individuals. The amount of activity-specific energy expended can then be estimated as energy $= 0.0175(\text{kcal/kg/min}) \times \text{weight (kg)} \times \text{MET}$ values [2]. Here, we use the ground truth labels to select activities to keep this procedure identical to the commonly used manual estimate. Table 8 presents the detailed results for each sequence. The accuracy is calculated over the total calorie expended in each recording session. We also measure the correlation between the ground truth and the observed values[*]. We can see that the fusion of visual and inertial sensors achieves higher accuracy and correlation in more sequences than the MET model or unimodal approaches, and obtains better rates on average, which points towards an advantage of using visual-inertial setups for the task of calorific expenditure prediction.

5 Conclusion and Future Directions

We have presented a system for calorific expenditure estimation using data from two different modality sensors, a RGB-Depth camera sensor and wearable inertial sensors (accelerometers). We have demonstrated the effectiveness of the fusion approach through a comprehensive comparative study with single modality setups and widely used METs prediction. The proposed fusion system used pooled spatial and temporal pyramids of visual and accelerometer features, which subsequently are fed in both early and late fusion approaches. To test the methodology, we introduced the challenging *SPHERE_RGBD+Inertial_calorie* dataset, which covers a wide variety of home-based human activities. The proposed fusion method demonstrates its ability to outperform the METs estimation approach and the use of single modality sensors. The focus of the paper has been on presenting a system for estimating calorific expenditure from combined visual and accelerometer sensors, where the purpose of the study has been to show that the fusion of both modalities improves the estimates beyond the accuracy of single modality, and the proposed system outperforms manual metabolic

---

[*]*Note that the total calorie values for sequence 4, 5, 8, 11, 15 and 16 are relatively low due to shorter sequences.*
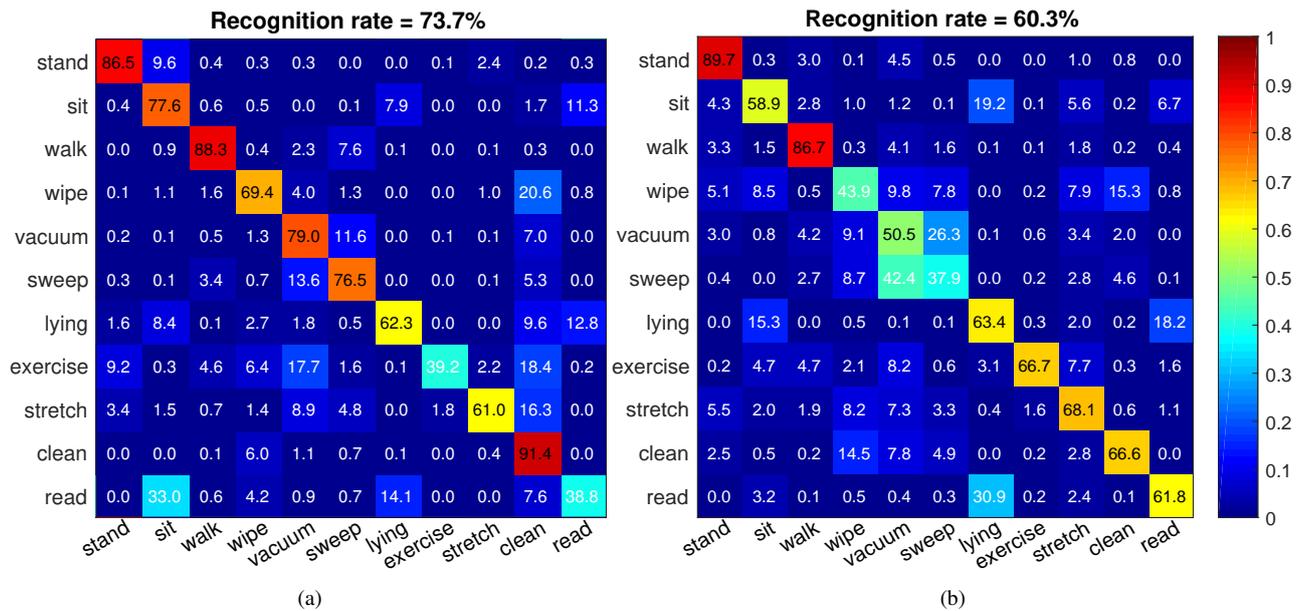
Fig. 6: The recognition confusion matrices from the best activity recognition results corresponding to the use of (a) visual sensor and (b) inertial sensors.

calorie w	activity w	*stand*	*sit*	*walk*	*wipe*	*vacuum*	*sweep*	*lying*	*exercise*	*stretch*	*clean*	*read*	**overall**
	n/a	**0.40**	0.45	**0.28**	**0.35**	0.32	0.38	**0.55**	**0.36**	0.55	0.36	0.50	**0.43**
	7.5	0.51	**0.43**	**0.28**	0.38	0.33	0.38	0.58	0.43	0.55	**0.35**	**0.43**	0.45
60	15	0.41	**0.43**	0.30	0.41	0.32	0.39	0.57	0.45	0.54	0.36	0.44	0.44
	30	0.47	0.44	0.30	0.42	**0.31**	**0.37**	0.56	0.44	0.53	0.37	0.46	0.44
	60	0.54	0.45	0.32	0.44	0.34	0.39	0.58	0.42	**0.52**	0.38	0.50	0.45

Table 5 Calorific expenditure prediction error (normalised RMSE) using the **visual sensor** when ground truth labels are used to select the activity-specific model (top row). and when action recognition is employed at different window lengths.

look up table based methods – the main measure used in clinical practice today. We acknowledge that applying more advanced fusion approaches and different feature representations may improve the performance further. Possible future directions include introducing deep learning models and investigating advanced data fusion methodologies for different modality sensors. We hope this work, and the new dataset, will establish a baseline for future research in the area.

6   Acknowledgement

7   References

1   G. Samitz, M. Egger, M. Zwahlen, Domains of physical activity and all-cause mortality: systematic review and dose–response meta-analysis of cohort studies, International journal of epidemiology 40 (5) (2011) 1382–1400.

2   B. Ainsworth, et al., Compendium of physical activities: an update of activity codes and met intensities, Medicine and science in sports and exercise 32 (9) (2000) 498–504.

3   E. Ravussin, S. Lillioja, T. Anderson, L. Christin, C. Bogardus, Determinants of 24-hour energy expenditure in man. methods and results using a respiratory chamber., Journal of Clinical Investigation 78 (6) (1986) 1568.

4   Cosmed K4b2, http://www.cosmed.com/.

5   C. Chen, R. Jafari, N. Kehtarnavaz, Improving human action recognition using fusion of depth camera and inertial sensors, IEEE Transactions on Human-Machine Systems 45 (1) (2015) 51–61.

6   H. Gjoreski, B. Kaluža, M. Gams, R. Milić, M. Luštrek, Context-based ensemble method for human energy expenditure estimation, Applied Soft Computing 37 (2015) 960–970.

7   J. Aggarwal, L. Xia, Human activity recognition from 3D data: A review, Pattern Recognition Letters 48 (2014) 70–80.

8   N. Zhu, T. Diethe, M. Camplani, L. Tao, A. Burrows, N. Twomey, D. Kaleshi, M. Mirmehdi, P. Flach, I. Craddock, Bridging e-health and the internet of things: The sphere project, IEEE Intelligent Systems 30 (4) (2015) 39–46.

9   P. Woznowski, X. Fafoutis, T. Song, S. Hannuna, M. Camplani, L. Tao, A. Paiement, E. Mellios, M. Haghighi, N. Zhu, et al., A multi-modal sensor infrastructure for healthcare in a residential environment, in: Communication Workshop (ICCW), 2015 IEEE International Conference on, IEEE, 2015, pp. 271–277.

10 R. Planinc, A. A. Chaaraoui, M. Kampel, F. Florez-Revuelta, Computer vision for active and assisted living, Active and Assisted Living: Technologies and Applications (2016) 57.

11 V. Pitsikalis, A. Katsamanis, S. Theodorakis, P. Maragos, Multimodal gesture recognition via multiple hypotheses rescoring, The Journal of Machine Learning Research 16 (1) (2015) 255–284.

12 S. Leutenegger, S. Lynen, M. Bosse, R. Siegwart, P. Furgale, Keyframe-based visual–inertial odometry using nonlinear optimization, The International Journal of Robotics Research 34 (3) (2015) 314–334.

13 W. Fang, L. Zheng, H. Deng, H. Zhang, Real-time motion tracking for mobile augmented/virtual reality using adaptive visual-inertial fusion, Sensors 17 (5) (2017) 1037.

14 S. Gasparrini, E. Cippitelli, E. Gambi, S. Spinsante, J. Wåhslén, I. Orhan, T. Lindh, Proposal and experimental evaluation of fall detection solution based on wearable and depth data fusion, ICT Innovations 2015 (2016) 99–108.

15 S. Stein, S. J. McKenna, Combining embedded accelerometers with computer vision for recognizing food preparation activities, Proceedings of the 2013 ACM international joint conference on Pervasive and ubiquitous computing (2013) 729–738.

16 T. Diethe, N. Twomey, M. Kull, P. Flach, I. Craddock, Probabilistic sensor fusion for ambient assisted living, arXiv preprint arXiv:1702.01209.

calorie w	activity w	stand	sit	walk	wipe	vacuum	sweep	lying	exercise	stretch	clean	read	overall
	n/a	0.47	**0.41**	**0.33**	**0.26**	**0.37**	0.45	**0.46**	**0.57**	0.47	0.40	**0.50**	**0.45**
	7.5	0.51	0.47	0.41	0.45	0.39	**0.40**	0.62	0.72	0.52	0.42	0.63	0.52
60	15	0.53	0.57	0.40	0.48	0.44	0.43	0.57	0.72	**0.46**	**0.39**	0.57	0.52
	30	0.48	0.45	0.46	0.50	0.43	0.43	0.57	0.68	0.55	0.41	0.58	0.53
	60	**0.41**	**0.41**	0.51	0.50	0.47	**0.40**	0.57	0.64	0.55	0.41	0.71	0.54

Table 6 Calorific expenditure prediction error (normalised RMSE) using the **inertial sensors** when ground truth labels are used to select the activity-specific model (top row), and when action recognition is employed at different window lengths.
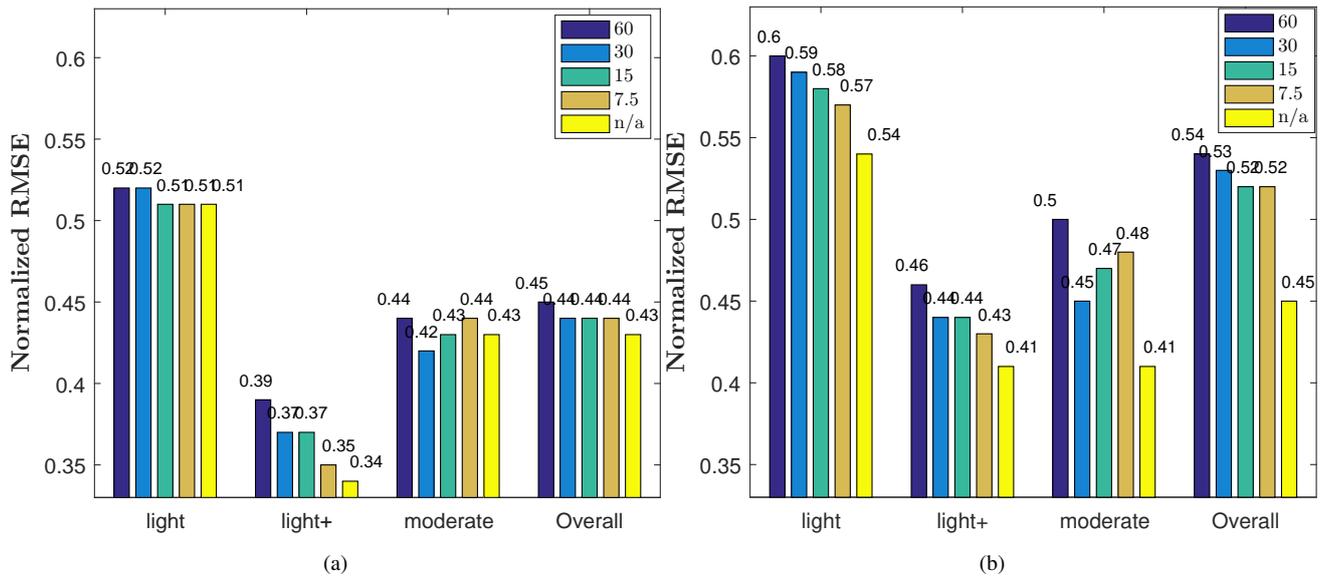
Fig. 7: **Prediction Accuracy of Calorific Expenditure.** Average calorie prediction errors (normalised RMSE) when ground truth labels are used to select the activity-specific model (in yellow), and when action recognition is employed at different window length using (a) a visual sensor and (b) two accelerometer sensors.

17 L. Tao, T. Burghardt, S. Hannuna, M. Camplani, A. Paiement, D. Damen, M. Mirmehdi, I. Craddock, A comparative home activity monitoring study using visual and inertial sensors, IEEE International Conference on E-Health Networking, Application and Services (2015) 644–647.

18 L. Tao, T. Burghardt, M. Mirmehdi, D. Damen, A. Cooper, M. Camplani, S. Hannuna, A. Paiement, I. Craddock, Real-time estimation of physical activity intensity for daily living, 2nd IET International Conference on Technologies for Active and Assisted Living (2016) 11–16.

19 L. Tao, T. Burghardt, M. Mirmehdi, D. Damen, A. Cooper, S. Hannuna, M. Camplani, A. Paiement, I. Craddock, Calorie counter: Rgb-depth visual estimation of energy expenditure at home, Asian Conference on Computer Vision, Workshop on Assistive Vision (2016) 0.

20 A. Edgcomb, F. Vahid, Estimating daily energy expenditure from video for assistive monitoring, International Conference on Healthcare Informatics (2013) 184–191.

21 P.-F. Tsou, C.-C. Wu, Estimation of calories consumption for aerobics using kinect based skeleton tracking, Systems, Man, and Cybernetics,International Conference on (2015) 1221–1226.

22 O. D. Lara, M. A. Labrador, A survey on human activity recognition using wearable sensors, IEEE Communications Surveys & Tutorials 15 (3) (2013) 1192–1209.

23 R. Igual, C. Medrano, I. Plaza, Challenges, issues and trends in fall detection systems, Biomedical engineering online 12 (1) (2013) 1.

24 I. Qudah, P. Leijdekkers, V. Gay, Using mobile phones to improve medication compliance and awareness for cardiac patients, Proceedings of the 3rd International Conference on PErvasive Technologies Related to Assistive Environments (2010) 36.

25 T. R. Bennett, J. Wu, N. Kehtarnavaz, R. Jafari, Inertial measurement unit-based wearable computers for assisted living applications: A signal processing perspective, IEEE Signal Processing Magazine 33 (2) (2016) 28–35.

26 N. Ravi, N. Dandekar, P. Mysore, M. L. Littman, Activity recognition from accelerometer data, AAAI 5 (2005) 1541–1546.

27 F. Ofli, R. Chaudhry, G. Kurillo, R. Vidal, R. Bajcsy, Berkeley mhad: A comprehensive multimodal human action database, Applications of Computer Vision (WACV), 2013 IEEE Workshop on (2013) 53–60.

28 D. Ravi, C. Wong, B. Lo, G.-Z. Yang, Deep learning for human activity recognition: A resource efficient implementation on low-power devices, in: Wearable and Implantable Body Sensor Networks (BSN), 2016 IEEE 13th International Conference on, IEEE, 2016, pp. 71–76.

29 D. Hendelman, K. Miller, C. Baggett, E. Debold, P. Freedson, Validity of accelerometry for the assessment of moderate intensity physical activity in the field., Medicine and science in sports and exercise 32 (9 Suppl) (2000) S442–9.

30 S. E. Crouter, J. R. Churilla, D. R. Bassett Jr, Estimating energy expenditure using accelerometers, European journal of applied physiology 98 (6) (2006) 601–612.

31 V. T. Hees, R. C. Lummel, K. R. Westerterp, Estimating activity-related energy expenditure under sedentary conditions using a tri-axial seismic accelerometer, Obesity 17 (6) (2009) 1287–1292.

32 M. Altini, J. Penders, O. Amft, Estimating oxygen uptake during nonsteady-state activities and transitions using wearable sensors, IEEE journal of biomedical and health informatics 20 (2) (2016) 469–475.

33 M. Altini, J. Penders, R. Vullers, O. Amft, Estimating energy expenditure using body-worn accelerometers: a comparison of methods, sensors number and positioning, IEEE Journal of Biomedical and Health Informatics 19 (1) (2015) 219–226.

34 J. Aggarwal, M. Ryoo, Human activity analysis: A review, ACM Computing Surveys 43 (3) (2011) 16.

35 M. Leo, G. Medioni, M. Trivedi, T. Kanade, G. Farinella, Computer vision for assistive technologies, Computer Vision and Image Understanding 154 (2017) 1–15.

36 G. Guo, A. Lai, A survey on still image based human action recognition, Pattern Recognition 47 (10) (2014) 3343–3361.

37 I. Laptev, On space-time interest points, International Journal of Computer Vision 64 (2-3) (2005) 107–123.

38 Y. Jia, E. Shelhamer, J. Donahue, S. Karayev, J. Long, R. Girshick, S. Guadarrama, T. Darrell, Caffe: Convolutional architecture for fast feature embedding, Proceedings of the ACM International Conference on Multimedia (2014) 675–678.

39 O. Oreifej, Z. Liu, Hon4d: Histogram of oriented 4D normals for activity recognition from depth sequences, Computer Vision and Pattern Recognition (2013) 716–723.

(a)

(b)
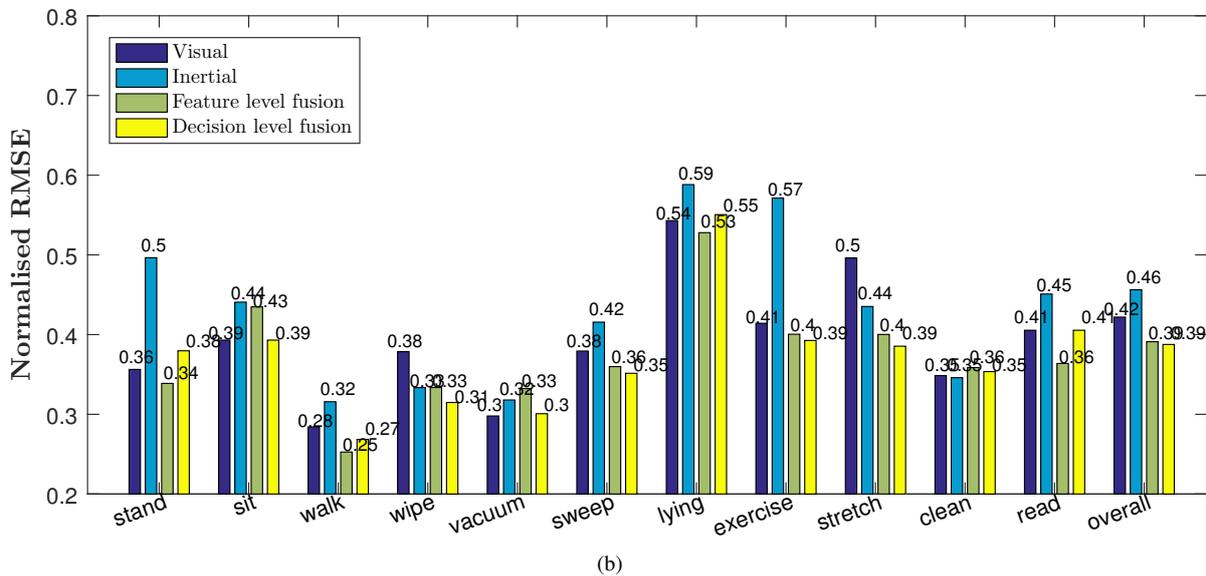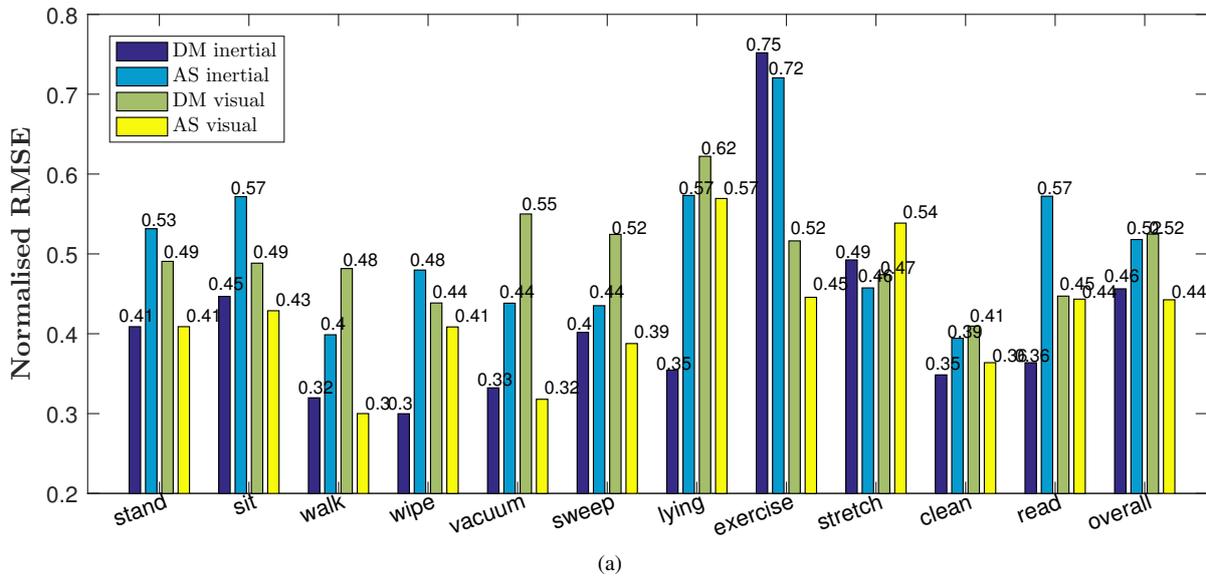
Fig. 8: (a) Average calorie prediction errors (normalised RMSE) of DM and AS approaches using visual and inertial sensors, respectively. (b) Average calorie prediction errors (normalised RMSE) of using visual sensor only (Visual), inertial sensors only (Inertial) and two sensor fusion approaches.

	w	*stand*	*sit*	*walk*	*wipe*	*vacuum*	*sweep*	*lying*	*exercise*	*stretch*	*clean*	*read*	**overall**
	Baseline	0.41	0.43	0.30	0.41	0.32	0.39	0.57	0.45	0.54	0.36	0.44	0.44
Visual	Recurrent1	0.36	**0.39**	**0.28**	**0.38**	**0.30**	**0.38**	**0.54**	0.41	0.50	**0.35**	**0.41**	**0.42**
	Recurrent2	**0.35**	0.49	0.31	0.52	0.59	0.53	0.56	**0.38**	**0.46**	0.42	0.45	0.52
	Baseline	**0.41**	**0.45**	**0.32**	**0.30**	**0.33**	**0.40**	**0.35**	**0.75**	**0.49**	**0.35**	**0.36**	**0.46**
Inertial	Recurrent1	**0.41**	0.54	0.37	0.32	0.34	0.41	0.45	0.78	**0.49**	0.36	0.44	0.49
	Recurrent2	0.50	1.06	0.63	0.49	0.45	0.49	0.92	0.95	0.52	0.44	0.85	0.68

Table 7 Average calorific expenditure prediction errors (normalised RMSE) for each activity with different learning approaches.

40 L. Tao, A. Paiement, D. Damen, M. Mirmehdi, S. Hannuna, M. Camplani, T. Burghardt, I. Craddock, A comparative study of pose representation and dynamics modelling for online motion quality assessment, Computer Vision and Image Understanding 148 (2016) 136–152.

41 I. Laptev, M. Marszałek, C. Schmid, B. Rozenfeld, Learning realistic human actions from movies, Computer Vision and Pattern Recognition (2008) 1–8.

42 F. Perronnin, J. Sánchez, T. Mensink, Improving the fisher kernel for large-scale image classification, European Conference on Computer Vision (2010) 143–156.

43 M. Ryoo, B. Rothrock, L. Matthies, Pooled motion features for first-person videos, Computer Vision and Pattern Recognition (2015) 896–904.

44 T. Dobhal, V. Shitole, G. Thomas, G. Navada, Human activity recognition using binary motion image and deep learning, Procedia Computer Science 58 (2015) 178–185.

45 L. L. Presti, M. La Cascia, 3d skeleton-based human action classification: A survey, Pattern Recognition 53 (2016) 130–147.

46 C. G. Snoek, M. Worring, A. W. Smeulders, Early versus late fusion in semantic video analysis, in: Proceedings of the 13th annual ACM international conference on Multimedia, ACM, 2005, pp. 399–402.

47 K. Liu, C. Chen, R. Jafari, N. Kehtarnavaz, Fusion of inertial and depth sensor data for robust hand gesture recognition, IEEE Sensors Journal 14 (6) (2014) 1898–1903.

48 J. Wu, J. Cheng, et al., Bayesian co-boosting for multi-modal gesture recognition., Journal of Machine Learning Research 15 (1) (2014) 3013–3036.

49 C. Chen, R. Jafari, N. Kehtarnavaz, A real-time human action recognition system using depth and inertial sensor fusion, IEEE Sensors Journal 16 (3) (2016) 773–781.

	Prediction (Calories)					Accuracy %				Correlation			
sequence	GT	Visual	Inertial	Fusion	MET	Visual	Inertial	Fusion	MET	Visual	Inertial	Fusion	MET
1	59	71	76	63	76	80.2	71.1	**93.9**	71.3	0.83	0.40	**0.88**	0.66
2	89	80	83	95	78	90.3	93.3	**93.5**	88.2	**0.85**	0.67	0.76	0.57
3	74	81	83	78	69	90.1	88.6	**94.1**	92.7	**0.84**	0.73	0.83	0.63
4	79	48	52	56	43	60.4	66.1	**71.5**	55.0	0.87	**0.89**	0.87	0.78
5	37	39	50	40	28	**98.6**	63.1	91.5	77.6	**0.90**	0.82	0.88	0.77
6	89	86	88	87	107	94.3	**99.0**	90.6	98.1	0.82	0.73	**0.83**	0.63
7	101	96	94	109	114	**95.3**	92.7	92.0	87.6	0.61	0.65	**0.66**	0.61
8	39	42	40	44	35	91.9	**97.2**	88.6	84.4	**0.93**	0.88	0.91	0.57
9	82	76	81	90	94	92.8	**99.3**	89.6	85.3	**0.86**	0.62	**0.86**	0.71
10	49	68	70	64	76	61.5	57.5	**68.6**	45.2	0.54	0.31	**0.55**	0.42
11	28	38	45	41	38	65.5	42.2	56.6	**66.8**	0.64	0.48	**0.66**	0.56
12	98	88	90	104	79	90.3	91.4	**93.8**	80.8	0.56	0.48	0.62	**0.66**
13	56	66	82	60	77	82.3	53.6	**91.9**	62.7	**0.78**	0.66	0.82	0.62
14	141	84	86	104	74	57.4	60.8	**73.8**	52.8	0.86	0.75	**0.87**	0.60
15	40	41	49	42	30	**98.9**	76.3	93.0	74.5	0.94	0.88	**0.96**	0.94
16	29	31	34	32	38	**97.3**	79.9	88.0	69.1	**0.88**	0.85	0.84	0.81
17	81	85	80	88	100	94.2	**99.5**	90.6	76.0	**0.74**	0.69	0.72	0.70
18	65	86	78	69	94	69.3	79.4	**93.1**	54.7	0.83	0.64	**0.85**	0.48
19	92	89	84	99	101	**94.2**	91.0	93.0	90.6	0.75	0.56	**0.77**	0.72
20	63	83	79	77	86	66.9	74.0	**77.0**	64.4	0.81	0.66	**0.86**	0.41
Average	-	-	-	-	-	82.9	78.8	**86.6**	73.7	0.80	0.67	**0.81**	0.64

Table 8 Ground truth and predicted calorie values in total per sequence and its accuracy and correlation. The best results for each sequence are in bold.

50   L. Breiman, Stacked regressions, Machine learning 24 (1) (1996) 49–64.
51   OpenNI organization, OpenNI User Guide (November 2010).
     URL http://www.openni.org/documentation
52   D. Tran, A. Sorokin, Human activity recognition with metric learning, European Conference on Computer Vision (2008) 548–561.
53   N. Dalal, B. Triggs, Histograms of oriented gradients for human detection, Computer Vision and Pattern Recognition 1 (2005) 886–893.
54   T. Dietterich, Machine learning for sequential data: A review, Structural, syntactic, and statistical pattern recognition (2002) 15–30.
55   C. V. Bouten, K. Koekkoek, M. Verduin, R. Kodde, J. D. Janssen, A triaxial accelerometer and portable data processing unit for the assessment of daily physical activity, Biomedical Engineering, IEEE Transactions on 44 (3) (1997) 136–147.
56   C. Chang, C. Lin, Libsvm: A library for support vector machines, ACM Transactions on Intelligent Systems and Technology 2 (3) (2011) 27.