

# Should I take a walk? Estimating Energy Expenditure from Video Data

Kunyu Peng\*, Alina Roitberg\*, Kailun Yang, Jiaming Zhang, and Rainer Stiefelhagen

\*denotes equal contribution

Karlsruhe Institute of Technology

firstname.lastname@kit.edu

## Abstract

We explore the problem of automatically inferring the amount of kilocalories used by human during physical activity from his/her video observation. To study this under-researched task, we introduce Vid2Burn – an omni-source benchmark for estimating caloric expenditure from video data featuring both, high- and low-intensity activities for which we derive energy expenditure annotations based on models established in medical literature.

In practice, a training set would only cover a certain amount of activity types, and it is important to validate, if the model indeed captures the essence of energy expenditure, (e.g., how many and which muscles are involved and how intense they work) instead of memorizing fixed values of specific activity categories seen during training. Ideally, the models should look beyond such category-specific biases and regress the caloric cost in videos depicting activity categories not explicitly present during training. With this property in mind, Vid2Burn is accompanied with a cross-category benchmark, where the task is to regress caloric expenditure for types of physical activities not present during training. An extensive evaluation of state-of-the-art approaches for video recognition modified for the energy expenditure estimation task demonstrates the difficulty of this problem, especially for new activity types at test-time, marking a new research direction. Dataset and code are available at <https://github.com/KPeng9510/Vid2Burn><sup>1</sup>.

## 1. Introduction

If you would ask people to honestly answer “Why do you go to the gym?” a frequent reply would be to burn calories. Physical activity is connected with our health and is an important element in prevention of obesity, diabetes or

<sup>1</sup>Acknowledgments: The research leading to these results was supported by the SmartAge project sponsored by the Carl Zeiss Stiftung (P2019-01-003; 2021-2026).

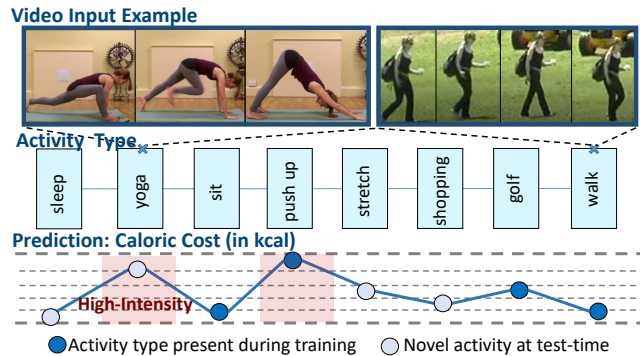


Figure 1. Our goal is to determine the amount of kilocalories burned during physical activity from video observation. To this intent, we collect Vid2Burn – an omni-source dataset with videos of people engaged in low- and high intensity activities and caloric cost annotations derived from established physiological models.

high blood pressure<sup>2</sup> – issues which are amplified through the recent Covid-19 lockdowns and the home office regulations [3]. With the rise of health tracking apps, automatic inference of energy expenditure is rapidly gaining attention [2, 5, 36, 39, 59, 70], but almost all prior research has focused on signals obtained from wearable devices, such as smart watches or heart rate monitoring chest straps. While such sensors are not always present at hand or comfortable to wear, most people can easily access a video camera in their phone or laptop. Apart from helping the users interested in tracking their exercise and maintaining active lifestyle, recent studies in gerontology highlight the benefits of automatically tracking the level of physical activity in assistive smart homes in order to support the elderly [4, 22, 45].

As important as it is for our health, understanding physical activity offers new technical challenges in computer vision. Excellent progress has been made in the field of human activity recognition [6, 14, 44, 56, 65] with remarkable accuracies reported on datasets such as HMDB-51 [29] or Kinetics [6]. However, when facing our task of estimating

<sup>2</sup>World Health Organization (WHO) - Physical inactivity a leading cause of disease and disability: <https://www.who.int/news/item/04-04-2002-physical-inactivity-a-leading-cause-of-disease-and-disability-warns-who>

caloric expenditure from human observations, these methods will face two main obstacles. First, the cornerstone of past research lies in rather rigid *categorization* into predefined actions. These categories are often relatively coarse, (e.g., “football” vs. “jogging”), so that the scene context provides the network with an excellent shortcut to the decision, leaving the actual moving person behind [10, 62]. Our task however requires *fine-grained* understanding of human movement, as medical research [7] lists *which muscles are active* and *how hard they work* as the main drivers of energy expenditure (although a multitude of further factors influence this complex physiological process). A second key challenge is to encourage the model to capture the essence of energy expenditure instead of memorizing average values of specific activity categories seen during training. Deep neural networks are prone to learning shortcuts [10, 16, 19] and internally casting the calorie regression problem as an “easier” task of activity categorization which might be one of such potential shortcuts. Even if the annotations are continuous calorie values and not rigid categories, in practice, the training set can only cover a finite amount of activity types. Ideally, our model should not be bounded to category-specific biases and indeed learn the nature of activity-induced energy expenditure by, e.g., understanding the type and intensity of bodily movement produced by the skeletal muscles. When developing an energy expenditure benchmark, it is therefore critical to evaluate the results on types of physical activity not present during training.

In this paper, we explore the new research direction of inferring activity-induced caloric cost by observing the human in video as Fig. 1. To tackle the lack of public large-scale datasets, we introduce *Vid2Burn* - a new omni-source benchmark spanning 9789 video examples of people engaged in different activities with corresponding annotations designed based on models established in medical literature [27] from (1) current activity category (2) intensity of the skeleton movement and (3) heart rate measurements obtained for a subset of activities (household activities) in a complementary study. Videos in the dataset are chosen from four diverse activity recognition datasets [25, 29, 50, 53] originally from YouTube, movies or explicitly designed for recognition in household context.

Yet, a key challenge when applying energy expenditure models in practice arises from transferring the learned concepts to new activity types. To meet this requirement, *Vid2Burn* is equipped with a cross-category benchmark, where the caloric cost estimation models are evaluated against activity types not seen during training. In addition to potential mobile health applications, our dataset there fills the lack of benchmark studying concise recognition of body movement without relying on category-specific context biases. From the computer vision perspective, the key technical challenges of our benchmark are (1) *fine-grained*

Datasets	Vid2Burn-Diverse	Vid2Burn-ADL	Vid2Burn
Video origin	Youtube/ movie datasets	ADL datasets	Youtube/ Movie/ADL
#Clips	4260	5529	9789
#Activities	33	39	72
#Train/test categories	27/6	33/6	60/12
Unit	kcal/hour	kcal/hour	kcal/hour
Min	64	153	64
Max	961	449	961
Mean	373	276	318

Table 1. An overview of the main properties of *Vid2Burn* and its two versions (including statistics of the caloric annotations).

understanding of bodily movement and (2) *generalization* to previously unseen types of activities. Extensive experiments with multiple state-of-the-art approaches for video- and body pose based action recognition demonstrate the difficulty of our task using modern video classification architectures, highlighting the need for further research.

## 2. Related work

**Activity recognition in videos.** Human activity recognition often operates on body poses [9, 31, 33, 51, 67] or learns representations end-to-end directly from the video data using Convolutional Neural Networks (CNNs) [6, 17, 20, 52, 68]. CNN-based approaches often deal with the temporal dimension via 3D convolution [44, 56, 57, 61, 65] or follow the 2D+1D paradigm, chaining spatial 2D convolutions and subsequent 1D modules to aggregate the features temporally [14, 15, 24, 32, 60, 71]. Fueled by multiple publicly released large-scale activity recognition datasets collected from Youtube/Movies [25, 29, 53] or in home environments [50], the research of deep learning based activity recognition became a very active research field also explored in more targeted applications, e.g., in cooking [11, 46], sports [42], robotics [23, 49], and automated driving [35]-related tasks. More specialized activity recognition research also addressed topics such as uncertainty of video classification models [47, 54]. However, all the approaches focus on *categorization* into previously defined activity classes, while examining their feasibility for capturing complex physiological processes of the body, such as our calorie expenditure task, has been largely overlooked.

**Energy expenditure prediction.** Visual estimation of calorie values has been mainly investigated in food image analysis (i.e., tracking the amount of caloric intake) [34, 40, 48]. Energy *expenditure* induced by physical activity is mostly studied from an egocentric perspective featuring data from wearable sensors, such as accelerometers or heart rate monitors [2, 5, 18, 26, 37, 39, 41, 55, 64], with a recent survey provided in [70]. Only very few works address the *visual* predicted activity-related caloric expenditure [39, 55]. The only dataset collected for energy expenditure prediction by

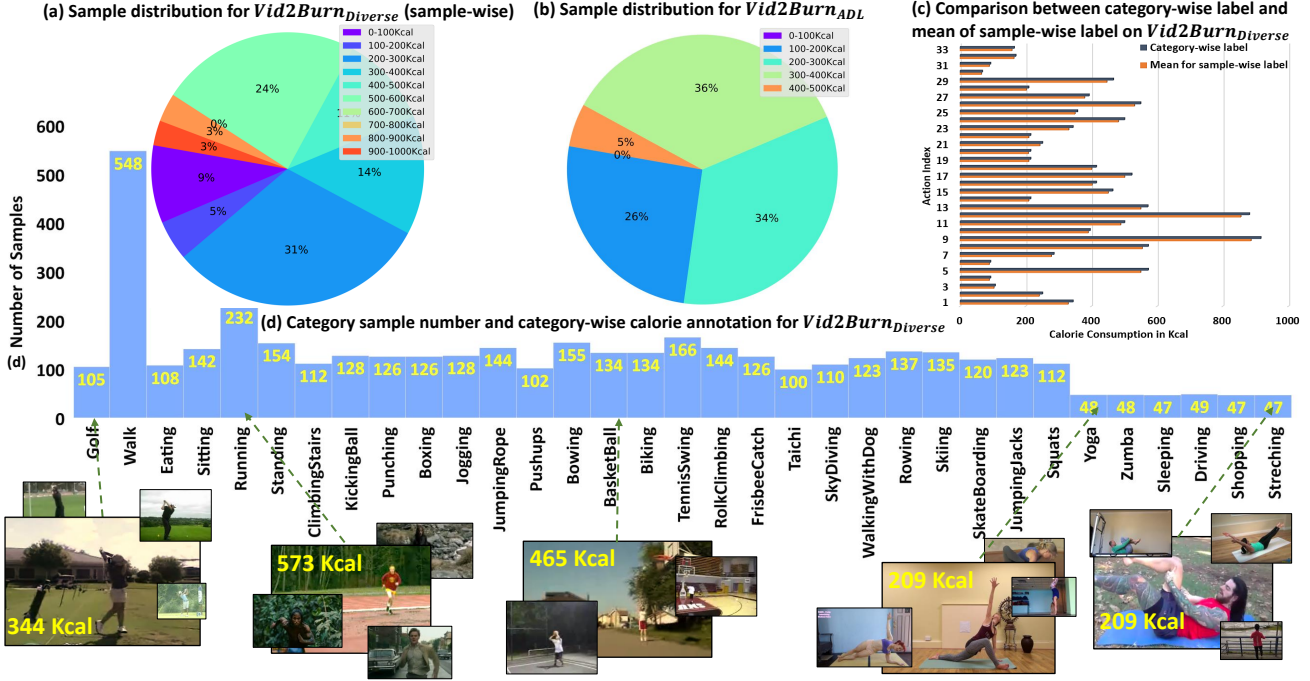


Figure 2. An overview of the dataset statistics. The statistics of the caloric cost values are summarized as pie charts in (a) and (b) for *Vid2Burn<sub>Diverse</sub>* and *Vid2Burn<sub>ADL</sub>* respectively. The caloric cost annotation statistics grouped by the individual activity categories are provided in (c), where blue bars represent the category-wise annotations, orange bars stand for the average of the sample-wise values and the bar index on the vertical axis indicates the action ID (which correspond to the order of activities listed below the blue histogram). The sample frequencies for different action categories in *Vid2Burn<sub>Diverse</sub>* are visualized in the blue bar chart (d) with multiple visual examples.

visually observing the human [55] features a highly simplistic evaluation setting (a single environment) and is comparably small in size, restricting the investigation of data-driven CNNs in this scenario (besides, the access to the collected database is restricted). To the best of our knowledge, no previous work explored deep CNNs for estimation of caloric cost from human observation. The research most similar to ours is presumably the work of Nakamura *et al.* [39], who collected an *egocentric* video dataset for estimating the energy expenditure and explore CNN-based architectures for this task. However, the research of [39] is significantly different from ours, as the cameras are mounted on the human, so that he/she is not observed. Our dataset is created with the opposite perspective in mind, as we target caloric expenditure estimation *from human video observations*.

### 3. Vid2Burn: A Benchmark for Estimating Caloric Expenditure in Videos

Given the growing demand for eHealth apps<sup>3</sup>, it is surprising that there is not a larger body of work on estimating physical intensity of activities in videos. This might be due

<sup>3</sup>Grand View Research. mHealth Apps Market Size, Share & Trends Analysis Report By Type (Fitness, Medical), By Region (North America, APAC, Europe, MEA, Latin America), And Segment Forecasts, 2021 - 2028. 2021. Available from: [www.grandviewresearch.com/industry-analysis/mhealth-app-market](http://www.grandviewresearch.com/industry-analysis/mhealth-app-market).

to the general focus of video classification research evolving mostly around activity *categorization* [6, 8, 13, 15, 17], while virtually all exercise intensity assessment datasets focus on *wearable* sensors [2, 5, 39] delivering, *e.g.*, heart rate or accelerometer signals. To promote the task of *visually* estimating the hourly amount of kilocalories burned by the human during the current activity, we introduce the novel *Vid2Burn* dataset, featuring > 9K videos of 72 different activity types with both caloric expenditure annotations on category- and sample-level.

#### 3.1. Dataset Collection

*Vid2Burn* is an omni-source dataset developed with a diverse range of movements and settings in mind. Our data collection procedure comprised the following steps. We started by surveying the well-known available datasets for categorical activity classification, (*e.g.*, [25, 29, 50, 53]). Then, we identified categories which are not only accessible from these public datasets but also have great technical feasibility to infer caloric cost annotations. The main sources of our dataset are UCF-101 [53], HMDB51 [29], test set of Kinetics [25] and NTU-RGBD [50]. We manually identified 72 activity types for which the hourly caloric cost can be estimated based on the established physiological models, (*e.g.*, [1, 27, 58]). Then, we estimated the labels for the energy expenditure based on these models on the

category- and sample-level described in Section 3.3.

### 3.2. Dataset Structure

The benefits of understanding caloric cost from videos extend to many applications, such as tracking of active exercise routines [21] or monitoring the daily physical activity level for elderly care [4, 22, 45]. From the technical perspective, it is also useful to distinguish settings with higher and lower differences between the samples. Lastly, while it is feasible to derive proper ground-truth for coarse behaviours or situations with well-studied energy expenditure, (e.g., types of sports and exercises), many daily living activities do not fall into this category and should be addressed with different techniques. Motivated by this, we group the content of *Vid2Burn* in two subsets: *Vid2Burn<sub>Diverse</sub>* and *Vid2Burn<sub>ADL</sub>*. Table 1 gives an overview of *Vid2Burn* and both its variants.

***Vid2Burn-Diverse*** is collected from Youtube- and movie-based sources [25, 29, 53] and therefore features a highly uncontrolled environment (camera movement, diverse inside/outside backgrounds). Since we focused on activities with well-studied energy expenditure models, a large portion of behaviours are related to sports, (e.g., *PushUps*). However, the database also covers certain everyday activities, such as *walking*, *standing* or *eating*. The distribution of different activity types is summarized in Figure 2. On average, the dataset features 129 video clips per category using category labels inherited from the original sources, with *walking* being unsurprisingly the most common behaviour for 548 videos while *stretching* and *shopping* are the least frequent ones for 47 videos.

***Vid2Burn-ADL*** on the other hand targets Activities of Daily Living (ADL) and might be used for physical workload tracking in smart homes. The activity types and video examples are derived from the public NTU-RGBD [50] dataset for ADL classification and, compared to *Vid2Burn<sub>Diverse</sub>*, this dataset contains activities of rather lower physical intensities (e.g., *pickup*, *take off jacket*, *read*, *drink water*). The environment of *Vid2Burn<sub>ADL</sub>* is much more controlled and the differences between the individual samples are at smaller scale. In other words, *Vid2Burn<sub>ADL</sub>* can be regarded as a much more fine-grained benchmark for caloric cost regression. In contrast to *Vid2Burn<sub>Diverse</sub>*, the categories of *Vid2Burn<sub>ADL</sub>* are rather well-balanced and the number of examples per activity type is 142 on average (detailed frequency statistics provided in the supplementary).

### 3.3. Caloric Expenditure Annotations

To adequately represent activity-induced caloric expenditure, we conducted a literature review on this physiological process [1, 37, 43, 58, 63]. Tracking the heat expended by nutrients oxidation (i.e., monitoring oxygen intake and carbon dioxide production) is considered the most accu-




Figure 3. Statistical analysis of the sample-wise caloric cost values in *Vid2Burn<sub>Diverse</sub>*. Red stars indicate the category-wise calorie consumption annotation, while the standard deviation boundaries result from sample-wise corrections based on the body movement.

rate way of estimating energy expenditure [43]. While this method is invasive and not practical for large-scale use, a multitude of topical studies conduct and publish such measurements for specific groups of activities, which is often summarized by meta-reviews in the form of compendiums [1]. Such catalogues provide energy expenditure values for specific activities that are often available online as look-up tables. A common way of estimating caloric cost, which is also leveraged by us, is deriving it from the heart-rate with validated physiological models [27, 37]. Our annotation scheme leverages three methods to estimate hourly energy expenditure: established medical compendiums [1], heart-rate based measurements [27] as well as adjustments based on the captured body movement [58].

Next, we describe the three different ways for obtaining caloric cost ground truth (Section 3.3.1) and explain how we leveraged them to annotate the *Vid2Burn<sub>Diverse</sub>* (Section 3.3.2) and the *Vid2Burn<sub>ADL</sub>* (Section 3.3.3) datasets.

#### 3.3.1 Annotation Methods

Our derived annotation scheme leverages three types of sources: (1) current activity category, (2) intensity of the skeleton movement, as well as, (3) heart rate measurements obtained for a subset of activities (household activities) in a complementary study.

**Caloric cost values from published compendiums.** First, we leverage activity-specific metabolic rate values from published compendiums [1] often summarized as look-up tables available on the web.<sup>1</sup> For simplicity, we assume the body weight as 150 lb, since this is also the average

<sup>1</sup><https://captaincalculator.com/health/calorie/>



body weight of subjects captured in our heart rate measurement study. Examples of the category-wise caloric expenditure annotations are marked as stars in Figure 3 for the *Vid2Burn<sub>Diverse</sub>* dataset.

**Heart-rate based annotations.** *Vid2Burn<sub>ADL</sub>* focuses on daily living activities, which naturally exhibit lower intensity of movement. The differences are at a much smaller scale compared to *Vid2Burn<sub>Diverse</sub>* and the average expected energy expenditure has not been well-studied for many such concise types of physical activity. However, due to the more restricted nature of the environment, the NTU-RGBD setting is easy to reproduce. In such cases, we recreate the environment of 39 activities of *Vid2Burn<sub>ADL</sub>* and estimate their average caloric cost based on *heart rate measurements* captured in a study with four volunteer participants. Four people, one female and three males, participated in the data collection (1 female, 3 male, 27.75 years old on average, average weight 150 lb). The participation in our study was voluntary, and the subjects were instructed about the scope and purpose of the data collection and have given their written consent according to the requirements of our institution.

The heart rate of all participants was recorded using a wrist band activity tracker.<sup>2</sup> The subjects were asked to execute 39 activities of the *Vid2Burn<sub>ADL</sub>* with a resting period in between to ensure the heart rate recovery. More information about the study setup is provided in the supplementary.

Given the measured heart rate, we compute the caloric cost of the activity in accordance to [27] as,

$$Cal_M = 60 \times T \times ((-55.0969) + (0.6309 \times HR) + (0.1988 \times W) + (0.2017 \times A))/4.184, \quad (1)$$

$$Cal_F = 60 \times T \times ((-20.4022) + (0.4472 \times HR) - (0.1263 \times W) + (0.074 \times A))/4.184, \quad (2)$$

where  $Cal_M/Cal_F$  indicates the hourly caloric expenditure for male/female,  $HR$  is heart rate,  $W$  is the body weight,  $A$  indicates participant's age and  $T$  is the time length in hour.

**Body movement- based annotations.** Next, we approximate the caloric cost induced by the movement by leveraging the model of Tsou *et al.* [58]. People can engage in the same type of activities in different ways and, since the amount of calories burned is directly linked to the amount/types of active muscles and the intensity, more active bodily movements lead to higher caloric cost. Tsou *et al.* [58] formalizes and validates a model corresponding based on the movement of eight body regions  $r$ . We estimate the skeleton movement using *AlphaPose* [12, 30, 66] for *Vid2Burn<sub>Diverse</sub>*, while for *Vid2Burn<sub>ADL</sub>* we use the skeleton data provided by the authors of the original datasets [38, 50]. Following [58], we group skeleton

joints into eight regions-of-interest to approximate the energy consumption as:

$$E_{body} = \sum_{t=1}^F \sum_{r=1}^8 \omega_r \lambda^2 \left| \frac{1}{2} M_r \Delta x_{rt}^2 + \frac{1}{2} M_r \Delta y_{rt}^2 + \frac{1}{2} M_r \Delta z_{rt}^2 \right| \quad (3)$$

where  $\Delta x_{rt}, \Delta y_{rt}, \Delta z_{rt}$  indicate the position difference of the body region between frame  $t$  and  $t - 1$  (the average position of all body joints inside same region),  $F$  indicates the frame number and  $\lambda$  is the frame frequency,  $M_r$  indicates the mass of  $r$ th body region. The final calorie consumption is obtained via multiplication with  $0.239(Cal/J) \times F \times \lambda \times 3600(h/s)$  in per hour-wise expression. The weighting factors  $\omega_i$  of the different body regions are obtained from [58]. The main purpose of the body-pose based caloric cost estimation is to enable more concise annotations at *sample-level*, since the same activity can be executed with different intensities.

### 3.3.2 Vid2Burn-Diverse annotations

One strategy behind the design of *Vid2Burn<sub>Diverse</sub>* was to select behaviour types for which the average caloric costs have been well-studied and easily accessible [1] (for that reason, sports-related videos constitute a significant portion of *Vid2Burn<sub>Diverse</sub>*). The main source for the category-level values is therefore derived from the published average category-specific which have been well-studied and easily accessible in this regard (on category level). We then correct the estimations of the individual videos based on the previously explained body-movement model [58], resulting in more concise sample-level annotations.

### 3.3.3 Vid2Burn-ADL annotations

Energy expenditure is not well-studied for many of the more concise daily living situations in *Vid2Burn<sub>ADL</sub>*. We therefore take a detour by conducting a study with participants' heart rate recorded during the 39 target activities (as described in the heart rate paragraph of Section 3.3.1). For each activity type, we estimate (1) the average heart rate-based caloric cost value obtained from our study, (2) the average values based on the skeleton movement captured in the *Vid2Burn<sub>ADL</sub>* videos and (3) if available, values established from published medical compendiums [1]. Note, that while for *Vid2Burn<sub>Diverse</sub>* such estimations from published compendiums represent the main ground truth source, which are only available for 4 out of 39 *Vid2Burn<sub>ADL</sub>* activities. The final category-level annotations are then computed as the average of the estimations delivered by the two/three aforementioned methods. Similar to *Vid2Burn<sub>Diverse</sub>*, we then derive more accurate sample-level annotations through previously described skeleton-based corrections.

<sup>2</sup><https://mi.com/global/mi-smart-band-4/specs/>

### 3.4. Same- and intra-activity splits

Since we specifically aim to rate generalization of the calorie estimation models for new activity types, we construct two testing scenarios covering (1) *Known activity types* evaluation, with videos covering the same behaviours as the training set and (2) *unknown activity types* evaluation where the train and test samples are drawn from different activity types. We randomly select 27 (*Vid2Burn<sub>Diverse</sub>*) and 33 (*Vid2Burn<sub>ADL</sub>*) activity types for the training set, while for both benchmark versions, the 6 remaining categories are used for evaluation. Next, the data of the 27/33 training activity types is further split into training/testing (with ratio 7 : 3) for the same-category evaluation. Note that the category annotations used for constructing the splits are inherited from the source datasets. Overall, our dataset comprises 2782/3243 videos for training, 1192/1390 samples for the validation on the same activity type and 286/896 samples for new-activity-type evaluation for the *Vid2Burn<sub>Diverse</sub>* and *Vid2Burn<sub>ADL</sub>* databases, respectively.

## 4. Activity Recognition Models in the Context of Caloric Cost Estimation

### 4.1. Learning continuous caloric values

Given a video input, our goal is to infer hourly energy cost of the activity in which the depicted human is involved. Note, that we target the *intensity* of the bodily activity and not its duration, *i.e.*, our goal is to infer kilocalories burnt *per hour*. Since our targets are *continuous* caloric values our task naturally suits regression-based losses, such as the Euclidean L2 loss. However, we observed that regression optimisation converges to a constant value in our case (a similar effect has been reported before in multimodal problems, *e.g.*, in [69]). We therefore address this problem as multinomial classification with additional label softening. Similar to [54], we binarize each caloric value annotation  $l$  with resolution of 1 kcal inside a range  $n \in [0, N]$ , where  $N$  is set to 1000 kcal. To keep certain regression properties, such as penalizing predictions which fall closer to the ground-truth bin less, we soften the labels through Gaussian distribution with a given standard deviation (STD) denoted as  $\sigma$ . Then, for each ground truth annotation  $l$ , we obtain the softened label as a distribution over  $N$  bins:

$$l_s[n] = \frac{1}{\sigma\sqrt{2\pi}} \exp\left(-\frac{(n-l)^2}{2\sigma^2}\right), \quad (4)$$

where  $l_s$  indicates the soft label used for the supervision. We then use the the Kullback-Leibler (KL) divergence between the ground truth and predicted distributions:

$$KL_{loss}(y, l_s) = \frac{1}{N} \sum_{i=0}^{N-1} (l_s[n](\log l_s[n]) - y[n]), \quad (5)$$

where  $y$  indicates the predicted distribution with  $n \in [0, N]$ .

Method	Known activity types			New activity types		
	MAE↓	SPC↑	NLL↓	MAE↓	SPC↑	NLL↓
Average	160.1	-	-	220.7	-	-
Random	306.4	-	-	342.8	-	-
ST-GCN	140.8	19.86	28.17	285.1	4.41	14.01
SF-AVR	34.5	30.46	5.23	159.5	8.20	11.50
R3D-AVR	49.5	29.30	5.01	154.3	8.56	14.58
I3D-AVR	38.9	32.25	4.70	228.7	7.61	16.32
R(2+1)D-AVR	47.4	30.58	4.63	246.9	10.52	14.26
SF-LSTM	39.7	31.71	4.85	227.3	10.45	17.41
R3D-LSTM	82.7	28.93	5.23	282.2	6.49	<b>12.16</b>
I3D-LSTM	57.9	31.41	4.87	251.9	8.77	12.26
R(2+1)D-LSTM	53.5	31.26	4.82	182.4	12.18	19.20

Table 2. Recognition results on the *Vid2Burn<sub>Diverse</sub>* benchmark in the *category-wise* caloric annotations setting.

Method	Known activity types				New activity types			
	Running		Climbing		Yoga		Shopping	
	SPC↑	MAE↓	SPC↑	MAE↓	SPC↑	MAE↓	SPC↑	MAE↓
ST-GCN	-13.51	310.9	14.81	225.6	19.65	230.5	-0.24	322.3
SF-AVR	11.76	43.8	24.19	57.7	20.18	101.9	3.43	174.4
R3D-AVR	6.14	102.3	21.13	135.5	17.82	97.4	1.17	108.5
I3D-AVR	7.69	82.6	13.96	81.8	13.10	191.9	1.92	190.5
R(2+1)D-AVR	12.88	288.0	25.03	74.0	15.20	217.9	4.68	271.1
SF-LSTM	12.59	9.6	15.79	86.7	19.49	133.6	3.91	266.8
R3D-LSTM	6.54	262.1	8.82	83.8	15.66	223.2	0.02	370.1
I3D-LSTM	6.09	147.7	14.85	99.9	12.04	217.6	4.15	280.0
R(2+1)D-LSTM	11.22	37.1	10.47	98.4	18.26	181.8	8.53	134.4

Table 3. Prediction quality for individual activity types on *Vid2Burn<sub>Diverse</sub>* (for two known and two new categories).

### 4.2. Video representation backbones

We adopt five modern video- and body pose-based architectures developed for categorical human activity recognition as our video representation backbones.

**I3D.** The Inflated 3D CNN (I3D) is a widely-used activity recognition backbone [6] and is a spatio-temporal version of the Inception-v1 network. Weights transfer from pretrained 2D CNNs and its pretraining is achieved by repeating (“inflating”) the weights along the temporal axis.

**R3D.** This 3D convolutional architecture [17] with a remarkable depth of 101 layers (enabled through residual connections) chains multiple ResNeXt blocks, which are shallow three-layered networks leveraging group convolution.

**R(2+1)D.** Unlike previous models, R(2+1)D [57] “mimics” spatio-temporal convolution by factorizing it into distinct 2D spatial and 1D temporal convolutions, yielding remarkable results despite these simpler operations. This framework also leverages a residual architecture.

**SlowFast.** Our last CNN-based architecture is the SlowFast model of Feichtenhofer *et al.* [13] which introduces two branches: a slow pathway and a fast pathway capturing cues from different temporal resolutions.

**ST-GCN.** In addition to the video-based models, we consider a popular architecture comprising a graph neural network operating on the estimated body poses [67], which uses spatial- and temporal graph convolutional neural network to harvest human motion cues.

### 4.3. Temporal fusion

Temporal windows captured by the above backbones vary between 16 frames for I3D [6], R(2+1) [57] and R3D [56] and 32 frames for SlowFast [15] are considerably smaller than the durations of the video clips captured in *Vid2Burn*. Given an input video of length  $T$  and a model  $f_\theta(\cdot)$  which takes as input  $F$  frames, we sequentially pass  $K$  video snippets  $\{\mathbf{t}_1, \mathbf{t}_2, \dots, \mathbf{t}_K\}$  using sliding window with overlapping resulting in  $K$  predictions  $\{f_\theta(\mathbf{t}_1), f_\theta(\mathbf{t}_2), \dots, f_\theta(\mathbf{t}_K)\}$ . We now consider two different strategies for fusing these results: (1) averaging of the output of the last fully connected layer and (2) learning to fuse the output with an additional LSTM network. The first method employs simple average pooling of the final representation:  $\text{pred}(\mathbf{t}) = \frac{1}{K} \sum_{i=1}^K f_\theta(\mathbf{t}_i)$ . Our second fusion strategy passes the last fully connected layer of our video representation backbone to an LSTM network with two layers with the number of neurons corresponding to the input size, trained together with the backbone model in an end-to-end fashion. Since an LSTM also produces sequential output, we also average the resultant sequence to obtain the final prediction.

## 5. Experiments

### 5.1. Evaluation Protocols

We adopt Mean Average Error (MAE) as our main evaluation metric and additionally report the Spearman Rank Correlation (SPC), and the Negative Logarithm Likelihood (NLL). MAE is an intuitive metric reporting the mean disparity between prediction and ground truth in our target units, (*i.e.*, kilocalories). Note, that while SPC illustrates the association strength between the ground truth and the predictions (it is +1 if one is a perfect monotone function of the other and -1 if they are fully opposed), it should be interpreted with care, since it ignores scaling and shifting of the data. In other words, SPC would not reflect if the number of kilocalories is constantly over/underestimated by similar amounts. It therefore should only be viewed as a complementary metric. Note, that we report SPC in % for better readability (*i.e.*, we multiply the result by 100). Our experiments are carried out in two annotation settings: category- and sample-wise annotations (see Section 3.3 for details). We view the sample-wise version as a more concise choice, but using static category labels is the protocol used in the past energy expenditure work from egocentric data [39] and is adopted for consistency. As explained in Section 3.4, we conduct the evaluations on both: behaviours present during training and new activity types.

### 5.2. Experiment Results

All implemented models greatly outperform the random and average baselines in all measures (Table 2), showing

Method	Known activity types			New activity types		
	MAE↓	SPC↑	NLL↓	MAE↓	SPC↑	NLL↓
Average	160.7	-	-	220.7	-	-
Random	308.2	-	-	341.1	-	-
SF-AVR	57.9	49.41	6.08	134.0	42.20	9.02
R3D-AVR	140.6	42.13	7.01	179.8	25.12	9.33
I3D-AVR	128.8	44.18	6.73	130.8	32.67	8.58
R(2+1)D-AVR	98.9	52.21	6.53	224.6	21.92	9.28

Table 4. Recognition results on the *Vid2BurnDiverse* benchmark in the *sample-wise* caloric annotations setting.

Method	#Params	Inference time	Same-MAE↓	New-MAE↓
SF-AVR	34.87M	102.44ms	34.5	159.5
R3D-AVR	34.04M	27.19ms	49.5	154.3
I3D-AVR	12.75M	65.22ms	38.9	228.7
R(2+1)D-AVR	33.87M	66.61ms	47.4	246.9
SF-LSTM	51.66M	102.73ms	39.7	227.3
R3D-LSTM	50.82M	27.30ms	82.7	282.2
I3D-LSTM	29.54M	65.75ms	57.9	251.9
R(2+1)D-LSTM	50.67M	66.82ms	53.5	182.4

Table 5. Nr. of parameters, inference time and MAE of different approaches (on *Vid2BurnDiverse*, Quadro-RTX 6000 graphic card).

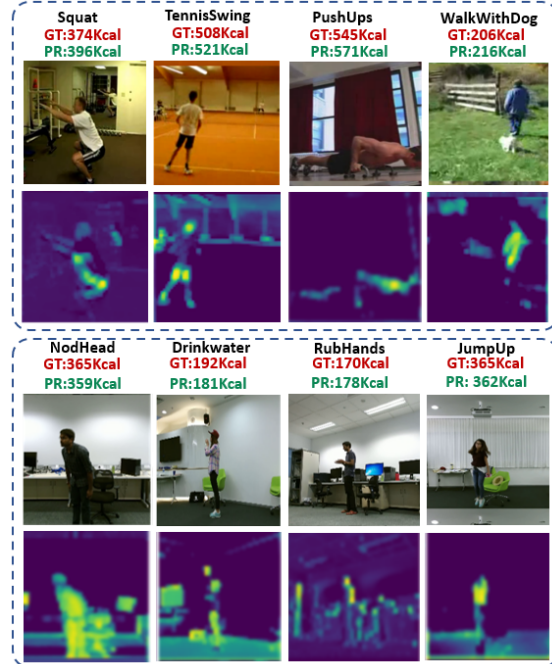


Figure 4. Qualitative results and visualization of the CNN activations for *Vid2BurnDiverse* (top) and *Vid2BurnADL* (bottom). We visualize the activations of the second convolutional layer (I3D). PR and GT indicate the prediction and the ground truth.

that our problem is feasible. However, our experiments also underline the difficulty of estimating the energy cost if the model is evaluated in new previously unseen situations, motivating further research of models with deeper and more fine-grained understanding of physical activities.

Tables 2 and 4 illustrate the performance of the employed algorithms as well as the random and average baselines for *Vid2BurnDiverse* in the category- and sample-labelled settings together with different temporal fusion schemes,

Method	Known activity types			New activity types		
	MAE ↓	SPC ↑	NLL ↓	MAE ↓	SPC ↑	NLL ↓
R3D-Reg	281.7	-	-	281.7	-	-
Average	76.2	-	-	59.4	-	-
Random	311.7	-	-	309.8	-	-
SF-AVR	20.1	68.61	5.57	36.4	72.30	6.69
R3D-AVR	35.1	76.91	6.01	44.3	79.78	6.21
I3D-AVR	22.9	72.97	5.66	39.6	70.45	6.38
R(2+1)D-AVR	26.2	71.31	5.99	29.5	75.30	6.28
SF-LSTM	36.9	77.19	6.05	47.1	80.43	6.24
R3D-LSTM	92.0	44.46	6.67	111.7	53.30	6.63
I3D-LSTM	101.4	42.42	6.68	68.1	51.11	6.66
R(2+1)D-LSTM	88.8	44.47	6.59	72.9	54.97	6.52

Table 6. Recognition results on the *Vid2Burn<sub>ADL</sub>* benchmark in the sample-wise caloric annotations setting.

marked as *AVR* (for the average pooling fusion) and *LSTM*, with *AVR* consistently leading notably better results. Slow-Fast (SF) with *AVG* consistently achieves the best recognition quality with only 34.5 kcal/20.1 kcal MAE. As expected, the results are lower in the more fine-grained sample-wise setting, since the backbones were initially developed for coarser categorization. The task of caloric cost regression in previously unseen situations is much more difficult and the performance drops significantly. The MAE is  $> 4$  and  $> 2$  times higher in the category- and sample-wise settings separately for SF-AVG. Interestingly, the best model in the case of *known activities* is usually not the top-performing approach for the *new activity types*, which is SF-AVG in case of the sample-wise evaluation of *Vid2Burn<sub>Diverse</sub>* with the MAE of 130.8 kcal, but the gap to SF-AVG (MAE of 134 kcal) is very small (Table 4).

Presumably due to a more restricted environment and smaller average caloric cost values, the models are more accurate on *Vid2Burn<sub>ADL</sub>* (Table 6). Consistently with the *Vid2Burn<sub>Diverse</sub>* results, SF-AVG yields the best recognition quality on known activities (MAE of 20.1 kcal) and is the second best performing model on the new ones. Table 6 also lists our results achieved with a regression loss (L2), marked as R3D-Reg. As explained in Section 4.1, we observe convergence to a constant value resulting in a very high MAE. Note that SPC cannot be computed in this case, since the output is a number while not a distribution. We also consider the trade-off between the performance and the computational cost (Table 5). *SF-AVG* offers a good balance between speed and accuracy, while the *I3D* backbone is a more lightweight model, but the MAE is higher.

We further look at the recognition for the individual activity types: two known (*running*, *climbing*) and two unknown (*yoga*, *shopping*) behaviours as presented in Table 3 (more results are provided in the supplementary). The recognition quality varies greatly depending on whether the activity is familiar. For example, SF-AVG is only off by 43.8 kcal for *running*, but the MAE is 174.4 kcal for *shopping*. Finally, in Figure 4 we showcase multiple qualitative results by visualizing the activation region of a CNN (I3D backbone) with multiple examples of representative qualitative results for the *Vid2Burn<sub>Diverse</sub>* (top) and *Vid2Burn<sub>ADL</sub>*

(bottom) datasets. Additionally to the predicted caloric value and the ground truth, we visualize the activation regions of an intermediate CNN layer (in this case we choose the second convolutional layer). It is evident, that the largest focus is put on the body region activated during movement which we view as a positive property, since the energy expenditure is a direct result of the muscle movement [7]. However, in several cases the objects (*e.g.*, the television) are clearly highlighted despite no direct interaction, indicating category-specific biases which are presumably the leading cause of mistakes in cross-category settings.

**Implementation Details.** Our models are trained with ADAM [28] using a weight decay of  $1e^{-5}$ , a batch size of 4, a learning rate of  $1e^{-4}$  for 40 epochs, the model weights from Kinetics [25] and a Quadro-RTX 6000 graphic card (parameter numbers and inference times reported in Table 5). For binarization of the continuous label space, the maximum caloric prediction limit is set to 1000 kcal for *Vid2Burn<sub>Diverse</sub>* and 500 kcal for *Vid2Burn<sub>ADL</sub>* with resolution of 1 kcal. A more detailed description of the parameter settings is provided in the supplementary.

## 6. Conclusion

We introduced the novel *Vid2Burn* benchmark for estimating the amount of calories burned during physical activities by visually observing the human. Through our experiments, we found that the generalization ability of modern video classification CNNs is limited in this challenging task and we will keep tackling this issue in the future work. *Vid2Burn* will be publicly released, opening new perspectives on specific challenges in human activity analysis, such as fine-grained understanding of bodily movement and generalization to new physical activity types, since our benchmark specifically evaluates the quality of energy expenditure estimation in *new* situations. We hope to foster research of human understanding models which are able to capture cues of the underlying physiological processes, (*e.g.*, active muscles and their intensity) instead of learning rigid category-specific biases seen during training.

**Broader Impact and Limitations.** This work targets energy expenditure estimation from videos. The benefits of such methods extend to multiple applications, such as supporting healthy lifestyle, *e.g.*, by tracking exercise routines [21] or monitoring the daily physical activity level for elderly care [45]. However, both annotations in our dataset and the results inferred by our models are *approximations* and not exact measurements, which should be carefully used in medical care applications, as they are simplified by assuming the body weight of 150 *lb*, while gender, height and age are not taken into account. Furthermore, our data-driven algorithms may learn shortcuts and biases present in the data potentially resulting in a false sense of security.



## References

- [1] Barbara E. Ainsworth, William L. Haskell, Stephen D. Herrmann, Nathanael Meckes, David R. Bassett, Catrine Tudor-Locke, Jennifer L. Greer, Jesse Vezina, Melicia C. Whitt-Glover, and Arthur S. Leon. 2011 compendium of physical activities: a second update of codes and MET values. *Medicine and Science in Sports and Exercise*, 2011. [3](#), [4](#), [5](#)
- [2] Fahd Albinali, Stephen Intille, William Haskell, and Mary Rosenberger. Using wearable activity type detection to improve physical activity energy expenditure estimation. In *UbiComp*, 2010. [1](#), [2](#), [3](#)
- [3] H. Amini, S. Habibi, A. H. Islamoglu, E. Isanejad, C. Uz, and H. Daniyari. COVID-19 pandemic-induced physical inactivity: the necessity of updating the global action plan on physical activity 2018-2030. *Environmental Health and Preventive Medicine*, 2021. [1](#)
- [4] Muhammad Awais, Lorenzo Chiari, Espen Alexander F. Ihlen, Jorunn L. Helbostad, and Luca Palmerini. Physical activity classification for elderly people in free-living conditions. *IEEE Journal of Biomedical and Health Informatics*, 2018. [1](#), [4](#)
- [5] Onur Barut, Li Zhou, and Yan Luo. Multitask LSTM model for human activity recognition and intensity estimation using wearable sensor data. *IEEE Internet of Things Journal*, 2020. [1](#), [2](#), [3](#)
- [6] Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? A new model and the kinetics dataset. In *CVPR*, 2017. [1](#), [2](#), [3](#), [6](#), [7](#)
- [7] Carl J. Caspersen, Kenneth E. Powell, and Gregory M. Christenson. Physical activity, exercise, and physical fitness: definitions and distinctions for health-related research. *Public Health Reports*, 1985. [2](#), [8](#)
- [8] Chun-Fu Richard Chen, Rameswar Panda, Kandan Ramakrishnan, Rogerio Feris, John Cohn, Aude Oliva, and Quanfu Fan. Deep analysis of CNN-based spatio-temporal representations for action recognition. In *CVPR*, 2021. [3](#)
- [9] Ke Cheng, Yifan Zhang, Xiangyu He, Weihan Chen, Jian Cheng, and Hanqing Lu. Skeleton-based action recognition with shift graph convolutional network. In *CVPR*, 2020. [2](#)
- [10] Jinwoo Choi, Chen Gao, Joseph C. E. Messou, and Jia-Bin Huang. Why can't I dance in the mall? learning to mitigate scene bias in action recognition. In *NeurIPS*, 2019. [2](#)
- [11] Dima Damen, Hazel Doughty, Giovanni Maria Farinella, Sanja Fidler, Antonino Furnari, Evangelos Kazakos, Davide Moltisanti, Jonathan Munro, Toby Perrett, Will Price, and Michael Wray. Scaling egocentric vision: The EPIC-KITCHENS dataset. In *ECCV*, 2018. [2](#)
- [12] Hao-Shu Fang, Shuqin Xie, Yu-Wing Tai, and Cewu Lu. RMPE: Regional multi-person pose estimation. In *ICCV*, 2017. [5](#)
- [13] Christoph Feichtenhofer, Haoqi Fan, Jitendra Malik, and Kaiming He. SlowFast networks for video recognition. In *ICCV*, 2019. [3](#), [6](#)
- [14] Christoph Feichtenhofer, Axel Pinz, and Richard P. Wildes. Spatiotemporal multiplier networks for video action recognition. In *CVPR*, 2017. [1](#), [2](#)
- [15] Christoph Feichtenhofer, Axel Pinz, and Andrew Zisserman. Convolutional two-stream network fusion for video action recognition. In *CVPR*, 2016. [2](#), [3](#), [7](#)
- [16] Robert Geirhos, Jörn-Henrik Jacobsen, Claudio Michaelis, Richard S. Zemel, Wieland Brendel, Matthias Bethge, and Felix A. Wichmann. Shortcut learning in deep neural networks. *Nature Machine Intelligence*, 2020. [2](#)
- [17] Kensho Hara, Hirokatsu Kataoka, and Yutaka Satoh. Learning spatio-temporal features with 3D residual networks for action recognition. In *ICCVW*, 2017. [2](#), [3](#), [6](#)
- [18] Mathias Hedegaard, Amjad Anvari-Moghaddam, Bjørn K. Jensen, Cecilie B. Jensen, Mads K. Pedersen, and Afshin Samani. Prediction of energy expenditure during activities of daily living by a wearable set of inertial sensors. *Medical Engineering & Physics*, 2020. [2](#)
- [19] Lisa Anne Hendricks, Kaylee Burns, Kate Saenko, Trevor Darrell, and Anna Rohrbach. Women also snowboard: Overcoming bias in captioning models. In *ECCV*, 2018. [2](#)
- [20] Shuiwang Ji, Wei Xu, Ming Yang, and Kai Yu. 3D convolutional neural networks for human action recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2013. [2](#)
- [21] Ara Jo, Bryan D. Coronel, Courtney E. Coakes, and Arch G. Mainous III. Is there a benefit to patients using wearable devices such as fitbit or health apps on mobiles? A systematic review. *The American Journal of Medicine*, 2019. [4](#), [8](#)
- [22] Tae Hee Jo, Jae Hoon Ma, and Seung Hyun Cha. Elderly perception on the internet of things-based integrated smart-home system. *Sensors*, 2021. [1](#), [4](#)
- [23] Michael Karg and Alexandra Kirsch. A human morning routine dataset. In *AAMAS*, 2014. [2](#)
- [24] Andrej Karpathy, George Toderici, Sanketh Shetty, Thomas Leung, Rahul Sukthankar, and Li Fei-Fei. Large-scale video classification with convolutional neural networks. In *CVPR*, 2014. [2](#)
- [25] Will Kay, João Carreira, Karen Simonyan, Brian Zhang, Chloe Hillier, Sudheendra Vijayanarasimhan, Fabio Viola, Tim Green, Trevor Back, Paul Natsev, Mustafa Suleyman, and Andrew Zisserman. The kinetics human action video dataset. *arXiv preprint arXiv:1705.06950*, 2017. [2](#), [3](#), [4](#), [8](#)
- [26] Bradley Kendall, Bryanne Bellovary, and Neha P. Gothe. Validity of wearable activity monitors for tracking steps and estimating energy expenditure during a graded maximal treadmill test. *Journal of Sports Sciences*, 2019. [2](#)
- [27] L. R. Keytel, J. H. Goedecke, Timothy D. Noakes, H. Hiloskorpi, Raija Laukkanen, Lize van der Merwe, and E. V. Lambert. Prediction of energy expenditure from heart rate monitoring during submaximal exercise. *Journal of Sports Sciences*, 2005. [2](#), [3](#), [4](#), [5](#)
- [28] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *ICLR*, 2015. [8](#)
- [29] Hildegard Kuehne, Hueihan Jhuang, Estíbaliz Garrote, Tomaso Poggio, and Thomas Serre. HMDB: A large video database for human motion recognition. In *ICCV*, 2011. [1](#), [2](#), [3](#), [4](#)
- [30] Jiefeng Li, Can Wang, Hao Zhu, Yihuan Mao, Hao-Shu Fang, and Cewu Lu. CrowdPose: Efficient crowded scenes

- pose estimation and a new benchmark. *arXiv preprint arXiv:1812.00324*, 2018. 5
- [31] Maosen Li, Siheng Chen, Xu Chen, Ya Zhang, Yanfeng Wang, and Qi Tian. Actional-structural graph convolutional networks for skeleton-based action recognition. In *CVPR*, 2019. 2
- [32] Ji Lin, Chuang Gan, and Song Han. TSM: Temporal shift module for efficient video understanding. In *ICCV*, 2019. 2
- [33] Ziyu Liu, Hongwen Zhang, Zhenghao Chen, Zhiyong Wang, and Wanli Ouyang. Disentangling and unifying graph convolutions for skeleton-based action recognition. In *CVPR*, 2020. 2
- [34] Javier Marin, Aritro Biswas, Ferda Ofli, Nicholas Hynes, Amaia Salvador, Yusuf Aytar, Ingmar Weber, and Antonio Torralba. Recipe1M+: A dataset for learning cross-modal embeddings for cooking recipes and food images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021. 2
- [35] Manuel Martin, Alina Roitberg, Monica Haurilet, Matthias Horne, Simon Reiß, Michael Voit, and Rainer Stiefelhausen. Drive&act: A multi-modal dataset for fine-grained driver behavior recognition in autonomous vehicles. In *ICCV*, 2019. 2
- [36] Alessandro Masullo, Tilo Burghardt, Dima Damen, Sion Hannuna, Víctor Ponce-López, and Majid Mirmehdi. CaloriNet: From silhouettes to calorie estimation in private environments. In *BMVC*, 2018. 1
- [37] Li Meina, Jang Myoung Kim, and Youn Tae Kim. A combined heart rate and movement index sensor for estimating the energy expenditure. In *SENSORS*, 2010. 2, 4
- [38] Alexander H. K. Montoye, Jordana Dahmen, Nigel Campbell, and Christopher P. Connolly. Accuracy of physical activity monitors for steps and calorie measurement during pregnancy walking. *Journal for the Measurement of Physical Behaviour*, 2019. 5
- [39] Katsuyuki Nakamura, Serena Yeung, Alexandre Alahi, and Li Fei-Fei. Jointly learning energy expenditures and activities using egocentric multimodal signals. In *CVPR*, 2017. 1, 2, 3, 7
- [40] Shu Naritomi and Keiji Yanai. CalorieCaptorGlass: Food calorie estimation based on actual size using HoloLens and deep learning. In *VRW*, 2020. 2
- [41] Ruairi O'Driscoll, Jake Turicchi, Mark Hopkins, Graham W. Horgan, Graham Finlayson, and James R. Stubbs. Improving energy expenditure estimates from wearable devices: A machine learning approach. *Journal of Sports Sciences*, 2020. 2
- [42] Paritosh Parmar and Brendan Tran Morris. What and how well you performed? A multitask learning approach to action quality assessment. In *CVPR*, 2019. 2
- [43] A. C. Pinheiro Volp, F. C. de Oliveira, R. Duarte Moreira Alves, E. A. Esteves, and Josefina Bressan. Energy expenditure: components and evaluation methods. *Nutricion Hospitalaria*, 2011. 4
- [44] Zhaofan Qiu, Ting Yao, and Tao Mei. Learning spatio-temporal representation with pseudo-3D residual networks. In *ICCV*, 2017. 1, 2
- [45] Blaine Reeder, Jane Chung, Kate Lyden, Joshua Winters, and Catherine M. Jankowski. Older women's perceptions of wearable and smart home activity sensors. *Informatics for Health and Social Care*, 2020. 1, 4, 8
- [46] Marcus Rohrbach, Sikandar Amin, Mykhaylo Andriluka, and Bernt Schiele. A database for fine grained activity detection of cooking activities. In *CVPR*, 2012. 2
- [47] Alina Roitberg, Monica Haurilet, Manuel Martinez, and Rainer Stiefelhausen. Uncertainty-sensitive activity recognition: A reliability benchmark and the CARING models. In *ICPR*, 2021. 2
- [48] Robin Ruede, Verena Heusser, Lukas Frank, Alina Roitberg, Monica Haurilet, and Rainer Stiefelhausen. Multi-task learning for calorie prediction on a novel large-scale recipe dataset enriched with nutritional information. In *ICPR*, 2021. 2
- [49] Lukas Rybok, Simon Friedberger, Uwe D. Hanebeck, and Rainer Stiefelhausen. The KIT robo-kitchen data set for the evaluation of view-based activity recognition systems. In *Humanoids*, 2011. 2
- [50] Amir Shahroudy, Jun Liu, Tian-Tsong Ng, and Gang Wang. NTU RGB+D: A large scale dataset for 3D human activity analysis. In *CVPR*, 2016. 2, 3, 4, 5
- [51] Lei Shi, Yifan Zhang, Jian Cheng, and Hanqing Lu. Skeleton-based action recognition with directed graph neural networks. In *CVPR*, 2019. 2
- [52] Karen Simonyan and Andrew Zisserman. Two-stream convolutional networks for action recognition in videos. In *NeurIPS*, 2014. 2
- [53] Khurram Soomro, Amir Roshan Zamir, and Mubarak Shah. UCF101: A dataset of 101 human actions classes from videos in the wild. *arXiv preprint arXiv:1212.0402*, 2012. 2, 3, 4
- [54] Yansong Tang, Zanlin Ni, Jiahuan Zhou, Danyang Zhang, Jiwen Lu, Ying Wu, and Jie Zhou. Uncertainty-aware score distribution learning for action quality assessment. In *CVPR*, 2020. 2, 6
- [55] Lili Tao, Tilo Burghardt, Majid Mirmehdi, Dima Damen, Ashley Cooper, Massimo Camplani, Sion Hannuna, Adeline Paiement, and Ian Craddock. Energy expenditure estimation using visual and inertial sensors. *IET Computer Vision*, 2018. 2, 3
- [56] Du Tran, Lubomir Bourdev, Rob Fergus, Lorenzo Torresani, and Manohar Paluri. Learning spatiotemporal features with 3D convolutional networks. In *ICCV*, 2015. 1, 2, 7
- [57] Du Tran, Heng Wang, Lorenzo Torresani, Jamie Ray, Yann LeCun, and Manohar Paluri. A closer look at spatiotemporal convolutions for action recognition. In *CVPR*, 2018. 2, 6, 7
- [58] Pei-Fu Tsou and Chao-Cheng Wu. Estimation of calories consumption for aerobics using kinect based skeleton tracking. In *SMC*, 2015. 3, 4, 5
- [59] Baodong Wang, Lili Tao, Tilo Burghardt, and Majid Mirmehdi. Caloric expenditure estimation using deep convolutional network features. In *WACVW*, 2018. 1
- [60] Limin Wang, Yuanjun Xiong, Zhe Wang, Yu Qiao, Dahua Lin, Xiaoou Tang, and Luc Van Gool. Temporal segment networks: Towards good practices for deep action recognition. In *ECCV*, 2016. 2

- [61] Xiaolong Wang, Ross Girshick, Abhinav Gupta, and Kaiming He. Non-local neural networks. In *CVPR*, 2018. 2
- [62] Philippe Weinzaepfel and Grégory Rogez. Mimetics: Towards understanding human actions out of context. *International Journal of Computer Vision*, 2021. 2
- [63] Robert T. Withers, Anthony G. Brooks, Simon M. Gunn, John L. Plummer, Christopher J. Gore, and John Cormack. Self-selected exercise intensity during household/garden activities and walking in 55 to 65-year-old females. *European Journal of Applied Physiology*, 2006. 4
- [64] Ling Xiao, Kai Wu, Xiaobing Tian, and Juan Luo. Activity-specific caloric expenditure estimation from kinetic energy harvesting in wearable devices. *Pervasive and Mobile Computing*, 2020. 2
- [65] Saining Xie, Chen Sun, Jonathan Huang, Zhuowen Tu, and Kevin Murphy. Rethinking spatiotemporal feature learning: Speed-accuracy trade-offs in video classification. In *ECCV*, 2018. 1, 2
- [66] Yuliang Xiu, Jiefeng Li, Haoyu Wang, Yinghong Fang, and Cewu Lu. Pose flow: Efficient online pose tracking. In *BMVC*, 2018. 5
- [67] Sijie Yan, Yuanjun Xiong, and Dahua Lin. Spatial temporal graph convolutional networks for skeleton-based action recognition. In *AAAI*, 2018. 2, 6
- [68] Joe Yue-Hei Ng, Matthew Hausknecht, Sudheendra Vijayanarasimhan, Oriol Vinyals, Rajat Monga, and George Toderici. Beyond short snippets: Deep networks for video classification. In *CVPR*, 2015. 2
- [69] Richard Zhang, Phillip Isola, and Alexei A. Efros. Colorful image colorization. In *ECCV*, 2016. 6
- [70] Xudong Zhang, Shuchen Cao, Hongbo Zhang, Jun Chen, M. M. Gupta, and Wenjun Zhang. Energy expenditure prediction methods: Review and new developments. In *IASC*, 2021. 1, 2
- [71] Bolei Zhou, Alex Andonian, Aude Oliva, and Antonio Torralba. Temporal relational reasoning in videos. In *ECCV*, 2018. 2