This item was submitted to [Loughborough's Research Repository](#) by the author.
Items in Figshare are protected by copyright, with all rights reserved, unless otherwise indicated.

A multimodal data processing system for LiDAR-based human activity recognition

PLEASE CITE THE PUBLISHED VERSION

<https://doi.org/10.1109/tcyb.2021.3085489>

PUBLISHER

Institute of Electrical and Electronics Engineers (IEEE)

VERSION

AM (Accepted Manuscript)

PUBLISHER STATEMENT

© 2021 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

LICENCE

All Rights Reserved

REPOSITORY RECORD

Roche, Jamie, Varuna De-Silva, Joosep Hook, Mirco Moencks, and Ahmet Kondo. 2021. "A Multimodal Data Processing System for Lidar-based Human Activity Recognition". Loughborough University.
<https://hdl.handle.net/2134/16592759.v1>.

A Multimodal Data Processing System for LiDAR Based Human Activity Recognition

Jamie Roche, Varuna De-Silva, Joosep Hook, Mirco Moencks & Ahmet Kondozi SMIEEE

Abstract—Increasingly, the task of detecting and recognising the actions of a human has been delegated to some form of neural network processing camera or wearable sensor data. Due to the degree to which the camera can be affected by lighting and wearable sensors scantiness, neither one modality can capture the required data to perform the task confidently. That being the case, range sensors like Light Detection and Ranging (LiDAR) can complement the process to perceive the environment more robustly. Most recently, researchers have been exploring ways to apply Convolutional Neural Networks to 3-Dimensional data. These methods typically rely on a single modality and cannot draw on information from complementing sensor streams to improve accuracy. This paper proposes a framework to tackle human activity recognition by leveraging the benefits of sensor fusion and multimodal machine learning. Given both RGB and point cloud data, our method describes the activities being performed by subjects using regions with a convolutional neural network (RCNN) and a 3-Dimensional modified Fisher vector network. Evaluated on a custom captured multimodal dataset demonstrates that the model outputs remarkably accurate human activity classification (90%). Furthermore, this framework can be used for sports analytics, understanding social behaviour, surveillance, and perhaps most notably by Autonomous Vehicles (AV) to data-driven decision-making policies in urban areas and indoor environments.

Index Terms—Convolution Neural Network, Fisher Vector, Faster RCNN, Human Activity Recognition, Multimodal Machine Learning.

I. INTRODUCTION

HUMAN activity recognition (HAR) has engaged many to become one of the most popular and active research fields in machine vision. Motivated by the academic value and practical applications in sports analytics [1], human-machine interface [2], ambient assisted living (AAL) [3] and autonomous vehicle (AV) research [4], [5], HAR is characterised by the actions taken by individuals in a group, or individuals acting alone. If a machine is to understand the actions taken by people, it needs to develop an intelligence that will allow it to reason about its behaviour.

To reason about the behaviour of people, three core challenges must be addressed. The primary challenge is getting a machine to understand the complex interaction performed by

Figure 1: Depicts the proposed HAR pipeline. Given RGB data, we first detect a subject and generate an ROI. Each RGB region of interest is translated and aligned onto the 3D point cloud. The corresponding 3D ROI is then passed onto the classifier before deciding what activity is being performed.

people when they modify their behaviour quickly depending on circumstances. These interactions are influenced by numerous factors, such as distance between agents or direction of travel. As such, the location and action of agents are equally important in facilitating a machine's ability to understand what people are doing.

The second challenge is the selection of attributes to be measured and the sensors that measure them. Being a popular and active research field, there are many different approaches. As such, sensor inaccuracy, constraints or obtrusiveness need to be considered. For example, wearable sensors provide reliable information about the actions of an agent. Unfortunately, wearable sensors are not as prolific as they need to be. While smartphone accelerometers are part of the way there, they cannot capture as much information as a camera can. Similarly, the camera is excellent at capturing a wide field of view but is prone to interference from lighting conditions. If each sensor has a shortcoming, the inevitable choice is to overcompensate on modality, thus improving the chances of capturing all the necessary data required to improve accuracy.

The third challenge is the sheer quantity of data required to train a classifier. Most contemporary supervised machine learning (ML) methods require examples to learn. In principle, given unlimited data, most supervised ML paradigms should be able to classify any activity under any circumstances. However,

there is a question of whether we can access the required data to achieve this task in practice. To that end, a framework that can utilise the data more intelligently is a significant challenge.

To address these challenges, we proposed a novel CNN architecture to leverage multimodal ML and sensing redundancy benefits. The proposed framework uses a pre-trained CNN (ResNet-50) [8] and a region proposal network (RPN) as an object detector [6]. The purpose of the object detector is to identify a region of interest (ROI) – containing a person performing an activity – in an image. The corresponding ROI is translated and aligned to point cloud data before being segmented and classified. This approach extends research on a 3D modified Fisher vector representation derived from a Gaussian Mixture Model (GMM) [7].

The proposed framework addresses the three challenges, showing that sensor data redundancy and multimodal ML enable us to use the data more intelligently, accommodate diverse feature space, and reason jointly about agents' activities and location. Specifically, we developed a technique where a detected ROI is projected onto LiDAR data, followed by a point cloud classification of an activity performed by a human, as depicted in Figure 1.

The primary contribution of this work is a new, accurate and reliable, method of HAR, which leverages multiple sensing modalities. The secondary and tertiary contributions are the Detection of an ROI and the Segmentation of subjects captured in 3D data. The rest of this paper is organised as follows: Section II provides an overview of the current state of the art in HAR. Section III details the proposed framework. Section IV reports on the comparative performance and visual results of the proposed framework. Section V discusses some of the limitations before concluding the paper in Section VI.

II. RELATED WORK

The technology and sensors integral to HAR vary to great a degree. This research uses a Camera and a LiDAR sensor to classify people performing different activities. The data collected by the camera is regular 2-Dimensional (2D) data, and the LiDAR collects 3D or point cloud data. To that end, we use 3D data and point cloud data interchangeably throughout the course of this article.

Multimodal ML is a research field with one of its earliest application in the field of Audio-Visual Speech Recognition (AVSR) [8], to contemporary fields of emotion recognition [9], [10] and tracking multiple agents [11] or Anomalies in traffic [12]. Like sensor data fusion, in the sense that it uses multiple sensor modalities, Multimodal ML uses multiple ML algorithms, either in parallel or in series, to improve a framework [13]. Five core technical challenges need to be addressed in Multimodal ML [14]. They can be described as; **representation**, which is concerned with the difference between data types; **translation** which is the mechanism of moving data from one plane to another; **alignment** which indicates the proximal relationship between two different sensor streams; **fusion** which denotes the joining of sensor data; and **co-learning** which means how knowledge is recognised from one modality for use with another.

Despite numerous developments in ML, the majority are demonstrated in tightly controlled settings with well-defined outcomes, typically perform singular tasks and generally use only one sensor stream [15]. Developing agents that make decisions using context information gathered from real-world environments remains an essential task, especially if we want a machine to understand what is going on around them.

The context of a scenario can be captured from various types of instruments, measurement techniques, and sensors. Also known as Multimodal sensing, it's where multiple sensor modalities capture context information about the environment they are working in [16], [17]. Commonly referred to as sensor data fusion, the challenges can be categorised into two groups [5], [16]: challenges at acquisition level and challenges due to the uncertainty of data sources.

One of the most prolific applications of Multimodal ML and sensor data fusion is in the area of AAL. Typically used in hospitals, rehabilitation centres and elderly care units [18], the algorithms are designed to determine what situations require urgent medical assistance. This process is usually done using multiple signals to determine a drift away from homeostasis [19] and identified using Multimodal ML by combining multiple biometrics to deduce what has just occurred.

A closely related research field is Multimodal ML for HAR and pedestrian detection [20]–[22]. Using the principles derived from AAL, it is expected that researchers could decrease road accidents and improve road safety. For example, authors in [23] and [24] provide an interesting application of Multimodal ML for a pedestrian recognition system that matches the pedestrians predicted intention with the driver's direction and, depending on the actions, brake the vehicle to avoid a collision.

For the most part, research into HAR has focused on either RGB, RGB-D or wearable sensor data. These techniques extract human silhouettes, tracking them in the temporal domain and classify activities based on patterns analysed [12] or velocity and acceleration data from an agent to make a classification. For example, research in [25]–[27] used microelectromechanical systems (MEMS) to facilitate a wearable sensor-based approach to activity recognition. Suitable for both indoor and outdoor environments, their size and lightweight make MEMS ideal for such scenarios. Furthermore, when combining multimodal ML methods in similar manners to researchers in [28], the argument for using wearable sensors becomes abundantly clear. Regrettably, though, MEMS have a long way to go to become ubiquitous. Even though smartphones are commonplace and can be used for this purpose, a single stream of low-level data, such as acceleration data, is insufficient to identify the more nuanced differences between walking with a phone and just walking [20], [29].

In some cases, a single stream of low-level data can be used to detect actions such as falls. However, in reality, the performance of such recognition systems is often significantly lower than the reported research results [30]. For example, authors in [31] cross-validated different fall detection algorithms and found unfavourable results. These unfavourable results can be attributed to variations in device hardware and

the relative position of sensors [32]. Clearly, when dealing with accelerometer data, the placement of sensors is paramount [33]–[38]. For example, falling, a desirable activity to recognise, and possible to do so with a single stream of low-level data from a smartphone, the type sensor and place where the user carries the device is a significant factor influencing the algorithms' ability to determine the action accurately. Although, possible to substitute a smartphone with wearable MEMS. We arrive back at the same problem of scantness, making the optical and range sensors an ideal substitute for the less than prolific wearable MEMS.

Wearable MEMS such as the Electrocardiogram (ECG) and Photoplethysmogram (PPG) can provide reasonably reliable information. However, particle applications like HAR suffer attenuation due to clothing artefacts [39], are easily influenced by diet [40] and provide erroneous data due to prescribed medication [41]. For example, artefactual signals caused by clothing can corrupt normal cardiac signal at the electrode. Dietary sodium, associated with unfavourable alterations in left ventricular mass, hypertension, and increased water retention, can significantly lower carotid-femoral pulse wave velocity (PWV). Conversely, medications such as L-Dopa, typically used in the treatment of Parkinson patients, causes postural hypotension – the pooling of blood in the lower limbs – resulting in an inability to distinguish the Diastolic peak of the PPG signal. This is not to say that wearable MEMS are not valuable sensors used to gather data, just that when applied to HAR they are scant and can easily be influenced by external factors, which result in inaccurate or erroneous data.

Herein the reason why some activity recognition researchers are focusing on RGB sensors rather than wearable MEMS. For example, authors in [1] propose a unified framework for HAR using a Recurrent Neural Network (RNN) to identify human subjects with an imaging sensor. Researchers in [42] used RNN's and imaging data to perform the same task except for single subject activity recognition. And researchers in [43] present a vision-based human activity recognition system using a mobile camera and histogram of orientated gradient (HOG) extracted from feature space.

Expanding on research using RGB, some scholars have made a further transition over to a data stream that can accommodate depth information and true colour. Using two parallel CNN's, called ResNet [44], to classify activities, authors in [45] presented a method of capturing the motion information of a whole video by producing a dynamic image corresponding to the input video. Similarly, authors in [46], [47] and [48] propose an activity recognition algorithm exploiting skeleton data extracted from RGB-D sensors, an activity scene recognition method based on a 3D skeleton sequence, and a probabilistic graphical model creating loopy subgraphs applicable for modelling simple and complex activities, respectively.

While most of the activity recognition systems have focused on wearable MEMS, RGB or RGB-D data, most contemporary systems are equipped with a camera and range sensors such as LiDAR. These sensors often complement traditional cameras (RGB-imaging sensors) to perceive the environment more robustly. Spurred on by events such as the 2004, 2005 defence

advanced research projects agency (DARPA) Grand Challenge and the 2007 DARPA Urban Challenge, the configuration of LiDAR and camera has become commonplace in almost all AV design [49],[50].

It is not just in AV research that LiDAR has been adopted as a data-gathering device. AAL research frequently uses LiDAR to make it easier for impaired people to navigate [51], [52]. Moreover, low-cost versions of the same LiDAR sensors are starting to enter the commercial world [53], [54], making them more accessible. This trend has led some to believe that LiDAR data, fused with image data, is the future of photography [55]. By removing the sensor's mechanical components, industry hopes to reduce the cost dramatically [27].

Concerning ML methods used for activity recognition, CNN's appear to be a popular choice [56]–[58]. Regrettably, point cloud data is not a natural input to CNN. CNN's were designed to process RGB images or a suitable data format, therefore adapting them to work on 3D data is not a straightforward extension. Depending on the sensor field of view (FoV), point cloud data is unstructured, unordered, prone to missing points, and affected by noise [7].

Early attempts at training CNN's using point cloud data required transforming the 3D data to a series of RGB images at multiple views. These networks learned the depth map of the scene rather than the 3D objects [59]. Although some success was attained, information was lost in the transformation, and the accuracy of the network suffered [60], [61].

More recently, a deep learning (DL) architecture called PointNet [62], [63] was proposed. During training, PointNet takes unordered data points as a set of functions and maps a point set onto a vector. When making a classification, the vector representation is checked against the patterns it identified in the dataset.

The advantage of PointNet is an end-to-end network that can process 3D data without translating it to a complex representation [63]. While PointNet is good at making a classification, it is void of a mechanism to facilitate detection and localisation. Moreover, the process is limited by available memory and lacks an understanding of the relationship between data points.

From the reviewed material, we understand the ML methods used for HAR vary between RNN's and CNN's, and one of three different types of sensors are used to capture data. In terms of practicalities, we know that wearable sensors can produce some exemplary results but are not suitable considering the infrastructure requirements and the ease at which data can be influenced. While RGB or RGB-D sensors capture more information, they are susceptible to lighting conditions and, therefore, not reliable.

Therefore, a solution for classifying human activities needs to rely on data that the sensors are providing, use it intelligently to rain in dimensionality and adequately capture all the necessary attributes without the need for ubiquitous infrastructure. This approach would have to borrow from multimodal ML techniques, sensor data fusion, and given the diverse environments humans interact with, the algorithm needs to be robust, with relative ease.

Figure 2: The proposed MfV Net. We first leverage the power of a CNN to detect objects and propose an ROI. The geometrically aligned and translated point cloud data allows us to identify the corresponding 3D ROI in the 3D sample. This segmented ROI is converted into a modified fisher vector representation before being passed onto the classification network. The resultant class is overlaid on the image, indicating the location and activity of the subject.

III. METHOD

In this section, we present our method of HAR using multimodal ML and sensing redundancy. The proposed method is based upon an advanced CNN architecture. The proposed framework illustrated in Figure 2 utilises a robust object detector to identify the ROI and a point cloud classifier to identify the subject's actions. The process requires the 2D ROI to be translated and aligned to the LiDAR data before the segmented point cloud data is classified into the activity being performed.

A. Object Detection

The principal component of the proposed architecture detects and localises agents before extracting the ROI from the RGB data. We modified ResNet [44] by training an RPN and a region classification network (RCN) to extract the ROI's. Referred to as a Faster-RCNN [6], this technique performs the convolution operation once per region before generating the feature map for the data. An extension of the RCNN and the Fast RCNN [6], the Faster RCNN's base classification network was ResNet-50. Trained on the ImageNet dataset [64] and finetuned using a subset (RGB images only) of the Loughborough London HAR (LboroLdnHAR) [3] dataset. This portion of the proposed framework was validated using an intersection-over union (IoU) between the ground truth bounding boxes and the predicted ROI. IoU detects the difference between the annotated bounding boxes of the RGB images and the proposed ROI. Using pairs of anchors with sizes $\{30,19;60,38;120,76\}$, the ROI was labelled positive (object of interest present) when accuracy was greater than 0.65 and negative (object of interest not present) when less than 0.35.

In this case, the accuracy the classifier returns for a region

behaves as a threshold for the detector. During training, the shortest side of all images was scaled to 246 pixels. Trained using stochastic gradient descent with a learning rate of 0.0001 and momentum of 0.9, the RCNN network used patch size of 16×16 pixels.

B. Geometric Alignment

The process of geometric alignment requires knowledge about proximity, orientation, and sensor modality. In terms of modality, there are vast differences between sensors that capture RGB, RGB-D and point cloud data. For example, point cloud data is a representation of a collection of 3D coordinates in space. Predominantly these coordinates can be captured in two ways: 3D scanning and photogrammetry [65]. The data is different in many ways to the structure and order of RGB data [61]. For example, the CCD or CMOS sensor at the heart of the camera captures reflected light from a subject. The specific location on the sensor that the reflected light strikes are recorded along with the true colour value of the data there capturing. Conversely, point cloud data captures the location of an object in 3D space. There is no sequence in which the ranges are recorded, and the scans vary in resolution, return either Spherical, Cylindrical or Cartesian coordinates of the objects they are sensing.

We address these issues by translating and aligning different sensor modalities. This process is performed in two stages. The foremost stage translates and aligns RGB-D data captured by the Kinect® sensor to the point cloud data captured by the LiDAR. The second stage translates the combined 3D data streams to the RGB camera data. It should be noted that in the first stage, the data types are comparable, so the process is relatively simple. In the latter, the process is more complicated

since we are translating different kinds of data. Figure 3 and Figure 4 illustrate a plan and elevated view of the sensor setup. This data is imperative to know so we can find the corresponding data points between different sensor modalities.

Figure 3: Illustrates the plan view of the setup showing the position of the sensors relative to the testbed.

For the process of geometric alignment, consider a point in space represented by a 3D coordinate. To shift that point's origin requires a scalar for each coordinate. For this explanation, consider an object O 4.76 m (x) in front of the Kinect® sensor, 2.75 m (y) to the right and 0.9 m (z) below the horizontal axis.

Figure 4: Illustrates the elevation view of the setup showing the position of the sensors relative to the testbed.

If we know that the Kinect® sensor is offset from the LiDAR – 0.45 m on the (Δx) axis, 0 m on the (Δy) axis and 0.5 on the (Δz) axis – it is a matter of adding the scalar to corresponding coordinates, to align the two data streams as per the set of Equation 1, 2, and 3.

$$X_T = \Delta x + x \quad (1)$$

$$Y_T = \Delta y + y \quad (2)$$

$$Z_T = \Delta z + z \quad (3)$$

In the second stage of the process, we find the corresponding 3D point for specific pixel output by the RGB sensor. This process is somewhat more complex than the foremost stage as the data types are not comparable.

For this derivation, consider an object O at a distance (c) of 5.6 m from the LiDAR. In this case, the object (O) is identified at an azimuth angle (B') of 3° and a zenith angle of 99° . The vertical distance (b) between the LiDAR and the RGB sensor is 27 cm, and the RGB sensor is positioned ($H_0 + \Delta z + b$) 122 cm above the ground. Considering the distance c between the object O and the LiDAR sensor, we can describe B , the relative

RGB azimuth angle as (4).

$$B = \sin^{-1} \left(\frac{\sin(A) \times b}{a} \right) \Rightarrow \sin^{-1} \left(\frac{\sin(99^\circ) \times 0.27}{5.62} \right) = 2.7^\circ \quad (4)$$

In this case, the RGB and LiDAR sensors are on the same longitudinal axis, and we know the LiDAR zenith angle A . As such, we can describe the RGB zenith angle C as (5):

$$C = 180 - A - B \Rightarrow 180 - 99 - 2.72 = 78.2^\circ \quad (5)$$

The distance a between object O and the RGB sensor can be described as (6):

$$a = \frac{\sin(C) \times a}{\sin(A)} \Rightarrow \frac{\sin(78.28) \times 5.62}{\sin(99^\circ)} = 5.5m \quad (6)$$

The purpose of this operation is to translate the LiDAR data stream to the image data. We assume that the longitudinal axis of the camera and the RGB sensor are aligned; however, an offset can be accounted for should it be required.

C. point cloud Classification

Unlike 2D RGB data, 3D data does not contain valuable descriptors – HOG, Hue Saturation and Variance (HSV) etc. – used by feature learning networks to classify the objects of interest. By its nature, 3D data needs to be modified before the CNN can classify it. The proposed method for classifying point cloud data consists of two main modules. The first module converts 3D data to the modified fisher vector representation. The second module processes the modified fisher vector representation in a CNN before making a classification. An in-depth description of the point cloud classifier used in this research can be found in [7].

1) Fisher Vector Representations

Intraclass variability describes variations in appearance between two views of the same class or variations in appearance between two instances of the same class. In activity recognition, for example, significant variations in appearance between two instances can occur. Fisher vectors are a way of dealing with intraclass variability and determine the likelihood of getting specific data given an underlying theory.

Figure 5: Shows three 2D points superimposed on a 2D Gaussian on the left of the image. In the centre of the image is the FV, and on the right is the FV representation for the three points. The vector in the centre of the image indicates the rate of curvature for the data.

A 2D representation of the fisher vector is shown in Figure 5. Formally, it is the expected value of the observed information. This can best be described as having a probability distribution with a very sharp peak in the ideal circumstance. Conversely, we can get a distribution where the curve is very

broad. This circumstance means there is a good deal of likelihood over a broad range of points.

Using a GMM, we can quantify the rate of curvature by looking at the probable distribution of the different points. It is often referred to as the curvature vector. It describes how curved the likelihood function is around the maximum. In simpler terms, the bigger the Fisher vector, the more curved the distribution is, meaning the more constrained the data is for that scenario.

Placing spherical Gaussians on a coarse grid provides structure, size, and the foundation for representing the image. If we compute each points derivative with respect to the Gaussian parameters and aggregate the results using three symmetric functions – maximum, minimum and summation – we get a fisher vector representation of a constant size that is invariant to permutations for different numbers of points.

For multiple points in a GMM, which consists of several Gaussian positioned on a grid, the Fisher vector representation can be interpreted as a statistically unique fingerprint of the 3D data [66]. Figure 6 (a) depicts the spherical Gaussians superimposed on the point cloud data, and Figure 6 (b) shows the modified Fisher vector representation. It should be noted that the modified Fisher vector representation in Figure 6 (b) is for visualisation purposes only.

Figure 6: (a) Shows the spherical Gaussians superimposed on the point cloud data. (b): Shows the modified Fisher Vector representation of the GMM for the point cloud data.

2) Classification of Fisher Vector Representation

The main parts of the network used to classify the representation comprise several inception modules, max-pooling layers, and four fully connected layers [67], [68]. Training is done using backpropagation and standard SoftMax and cross-entropy loss with batch normalisation. Dropout is placed on the fully connected layers to prevent excessive co-adapting and overfitting during training. Dropout refers to a technique of ignoring specific neurons response during a forward or backwards pass after a certain time. Between the last max-pooling layer and the first fully connected layer, the network has approximately 4.6 million trained parameters.

The overall objective of the inception module is to overcome the dimensionality of the multiplier effect. An Inception module

allows us to use a number of filters in a single module [68]. While this will make the network architecture wider and more complex, the process works remarkably well in view of what we are trying to achieve.

The penultimate element of the classification network is a max-pooling layer. The max-pooling layer allows us to reduce the dimensionality of the inception module further.

Max pooling layers work by passing all output elements of the inception module through an $n \times n$ filter. The $n \times n$ filter passes over the data with a fixed stride. Taking the first $n \times n$ region, we calculate the max value for that region before passing it on to the next. This process is repeated – passing the filter overall data elements – using the stride value to shift the pooling filter over in increments.

The final component in the classification network is four fully connected layers. The output of the previous layers is flattened into a single vector of values. These represent a probability that a particular feature belongs to a specific class. For example, in HAR, when a person is running, the fully connected layers should identify features representing the action with a high probability.

IV. EXPERIMENTAL RESULTS

This section is organised into six parts. The first section summarises the LboroLdnHAR dataset developed in [3]. The second, third and fourth sections report on the performance of the proposed framework, PointNet [62] and the Deep Neural Network framework (DNN) [3], [46], for HAR, respectively.

Comparative evaluation of the framework in this manner is essential to appraise the performance of the proposed process against state-of-the-art methods. It is not always clear how to do this, as what works for one framework will not necessarily work for another, particularly when they process different data types. With this in mind, we have explained the metrics we use and the reason for using them.

In the fifth section, we benchmark the proposed framework against baseline methods of HAR using the SYSU-3D dataset before presenting some visible results of the proposed framework in the final section.

A. Dataset

Existing techniques for HAR are sensitive to viewpoint variations. This is a consequence of extracting features from viewpoint dependent images. In contrast, we mine RGB data to identify a region in a point cloud scan. Since we segment and classify viewpoint independent data, this technique is insensitive to viewpoint variations. To achieve this, we require a specific type of dataset to fulfil the requirements of multimodality while providing optical and range data from at least two sensor streams of a known location.

The LboroLdnHAR [3] dataset consists of 6712-point cloud, RGB-D and RGB annotated by hand samples. Data was captured by three sensors, a VLP-16 LiDAR, Kinect® and 360° Camera. The dataset contains data of 9 typical human activities performed by 16 participants. Activities were chosen on the basis of realistic activities performed in an office environment. The activities were: carrying boxes, lying down, pushing a

board, running; sitting on a chair; sitting on a stool; standing while texting; walking and walking while texting.

The mean period for each data capture was approximately 35 seconds. Subjects started and ended the data capture periods with a T pose – standing upright with arms outstretched. During the experiments, the RGB 360° Camera, RGB-D Kinect® V2 and VLP-16 LiDAR sensor logged data at 30 fps, 30 fps and 5 fps, respectively.

We split the dataset into three subsets. The first subsets consisted of images with annotated ground truth ROI labels and trained the object detector. The second subsets contained aligned Kinect® and LiDAR data. This subset of data was used to train the point cloud classifier. The final subset was a small number of the aligned and translated point cloud, RGB-D and RGB samples not used during the training or validation of the proposed framework. This subsection of the dataset was used to demonstrate the performance of the proposed framework visually.

B. Proposed Framework Performance

In this section, we report on the performance of the individual components of the proposed framework. We chose to evaluate the proposed framework in this manner to assist in the selection of a point cloud classifier. Not all point cloud classifiers are suitable for activity recognition. Some do not retain the crucial relationships between different points that allow us to understand the system of rigid segments connected by articulated joints that enable humans to move. Similarly, when comparing the performance of the proposed framework, it would make little sense to compare against a framework that used wearable MEMS as the data is dramatically different. To that end, we chose to evaluate the individual components of the framework separately and compare the methods of classification against an alternative point cloud classifier and a DNN rather than a classifier that process MEMS data.

Figure 7: The average precision of the object detector portion of the proposed framework showing the trade-off between precision and recall.

1) Object Detection Performance

The location and detection of agents were performed using an object detector. We trained and evaluated the object detector on the first subset of the LboroLdnHAR dataset [3]. To measure the detector’s performance, we focused on the average

precision – shown in Figure 7 – the log-average miss rates – shown in Figure 8 – and the F-score detailed in Table 1. The average precision is a method of evaluation that incorporates the ability of the detector to make correct classifications, i.e., precision, and the ability of the detector to find all relevant objects, i.e., recall.

Figure 8: The log-average miss rate of the object detector portion of the proposed framework indicating the quality of detection of the object detector.

The log average miss rate was attained by varying the thresholds on the detector confidence prediction and measure the rate of change in the true and false positives. Generally speaking, the log-average miss rate returns the results of the object detector compared to the ground truth table. This metric is another method of measuring the performance of the object detector [69], [70]. In this case, we found that the average precision and log average miss rate for our detector was 0.95 and 0.1, respectively.

TABLE 1: OBJECT DETECTION PERFORMANCE

	PERFORMANCE
PRECISION	0.95
RECALL	0.96
LOG-AVERAGE MISS RATE	0.1
F-SCORE	0.95

The F-score considers both the precision and the recall of the detector. Using binary classification in conjunction with the F-Score allows us to determine the framework’s accuracy [71]. Defined as two times the precision times the recall divided by the precision plus the recall, it reaches its best value at 1 and worst at 0. We determined the F-Score – listed in Table 1 – as 0.95 for the object detector. These metrics are frequently used to determine the performance of an object detector

2) point cloud Classifier Performance

Table 2 reports the average score for each class when tested with the LboroLdnHAR dataset [3]. To measure the classifier’s performance, we focused on the precision, recall and the F-score of each class. Figure 9 shows the confusion matrix for the point cloud classifier. On the Y-axis are the Output Class, and on the X-axis is the Target Class. The diagonal cells indicate true positives that are correctly classified. The off-diagonal cells indicate false positives that are incorrectly classified. The overall accuracy for the classifier was 0.903. The Precision and Recall were 0.895 and 0.894, respectively.

		Confusion Matrix									
Output Class	CARRYING BOXES	71	0	0	0	0	0	0	0	0	100%
	LYING DOWN	0	103	0	0	0	0	0	0	0	100%
	PUSHING A BOARD	0	0	86	0	0	0	0	1	0	98.9%
	RUNNING	0	0	0	53	2	0	0	8	10	72.6%
	SITTING ON A CHAIR	3	0	0	5	89	3	0	0	1	88.1%
	SITTING ON A STOOL	0	0	0	0	1	91	1	0	1	96.8%
	STANDING WHILE TEXTING	0	0	0	0	8	1	95	0	0	91.3%
	WALKING	0	0	0	0	14	0	0	52	6	72.2%
	WALKING WHILE TEXTING	0	0	0	5	0	0	0	6	70	86.4%
		95.9%	100%	100%	68.8%	89.0%	95.8%	99.0%	77.6%	79.5%	90.3%
		4.1%	0.0%	0.0%	31.2%	11.0%	4.2%	1.0%	22.4%	20.5%	9.7%
		Target Class									
		CARRYING BOXES	LYING DOWN	PUSHING A BOARD	RUNNING	SITTING ON A CHAIR	SITTING ON A STOOL	STANDING WHILE TEXTING	WALKING	WALKING WHILE TEXTING	

Figure 9: Confusion Matrix for the classifier portion of the proposed framework.

C. PointNet Performance

We compared the performance of the proposed method against PointNet [62]. PointNet [62] was trained and tested with the LboroLdnHAR dataset [3]. Table 3 reports the average score for each class from PointNet. Similar to the previous section, where we focused on the precision, recall and F-score.

		Confusion Matrix									
Output Class	CARRYING BOXES	1	8	0	3	30	5	0	1	3	2.0%
	LYING DOWN	0	0	0	0	0	0	0	0	0	NaN%
	PUSHING A BOARD	0	0	5	0	2	0	0	0	0	71.4%
	RUNNING	67	21	77	67	35	76	96	62	84	11.5%
	SITTING ON A CHAIR	0	0	0	0	0	0	0	0	0	NaN%
	SITTING ON A STOOL	0	0	0	0	0	0	0	0	0	NaN%
	STANDING WHILE TEXTING	0	0	2	0	0	0	0	0	0	0.0%
	WALKING	2	73	1	7	3	14	0	4	1	3.8%
	WALKING WHILE TEXTING	3	1	2	0	30	0	0	0	0	0.0%
		1.4%	0.0%	5.7%	87.0%	0.0%	0.0%	0.0%	6.0%	0.0%	9.8%
		98.6%	100%	94.3%	13.0%	100%	100%	100%	94.0%	100%	90.2%
		Target Class									
		CARRYING BOXES	LYING DOWN	PUSHING A BOARD	RUNNING	SITTING ON A CHAIR	SITTING ON A STOOL	STANDING WHILE TEXTING	WALKING	WALKING WHILE TEXTING	

Figure 10: Confusion Matrix for PointNet. PointNet was tested on the same data used to test the proposed framework.

Figure 10 shows the confusion matrix for PointNet [62]. The overall accuracy for the classifier was 0.098. The Precision and Recall were 0.111 and 0.147, respectively. Of particular interest was the networks inability to classify subjects lying down, sitting on a chair, and sitting on a stool. While the network was able to make some predictions, its performance was substandard compared to the 3D Classification network.

		Confusion Matrix									
Output Class	CARRYING BOXES	21	0	0	0	0	0	0	0	10	67.7%
	LYING DOWN	0	103	0	0	0	0	0	0	0	100%
	PUSHING A BOARD	0	0	77	1	0	0	0	0	0	98.7%
	RUNNING	10	0	9	64	0	0	0	0	0	77.1%
	SITTING ON A CHAIR	0	0	0	2	101	0	0	0	0	98.1%
	SITTING ON A STOOL	0	0	0	0	0	94	0	0	0	100%
	STANDING WHILE TEXTING	0	0	0	1	0	0	104	0	0	99.0%
	WALKING	7	0	0	5	0	0	0	41	18	57.7%
	WALKING WHILE TEXTING	36	0	0	0	0	0	0	31	53	44.2%
		28.4%	100%	89.5%	87.7%	100%	100%	100%	56.9%	65.4%	83.5%
		71.6%	0.0%	10.5%	12.3%	0.0%	0.0%	0.0%	43.1%	34.6%	16.5%
		Target Class									
		CARRYING BOXES	LYING DOWN	PUSHING A BOARD	RUNNING	SITTING ON A CHAIR	SITTING ON A STOOL	STANDING WHILE TEXTING	WALKING	WALKING WHILE TEXTING	

Figure 11: Confusion Matrix for the DNN framework. The DNN framework was tested on the same data used to test the proposed framework.

D. DNN Performance

To make an additional comparison, we benchmark the performance of the proposed method against a DNN framework for HAR reported in [3]. This is a similar framework reported in [46] and a frequently used method to identify human activities from RGB-D data. The DNN framework was trained and tested with the LboroLdnHAR dataset [4]. Table 4 reports on the precision, recall and the F-score of each class from the DNN framework. Figure 11 shows the confusion matrix for the DNN framework. The overall accuracy for the classifier was 0.835. The Precision and Recall were 0.808 and 0.825, respectively. Of particular interest was the network's performance in classifying subject when moving. For example, the accuracy when classifying stationary subjects was remarkably higher than those from a dynamic class. This is evident in Figure 11, where the framework frequently misclassifies walking, running, and carrying boxes.

E. Performance on the SYSU-3D Dataset

To benchmark the proposed framework, we trained and tested the network on the joint data from the SYSU-3D dataset. The SYSU-3D dataset consists of 40 subjects performing 12 activities: drinking, pouring, calling phone, playing phone, wearing backpacks, packing backpacks, sitting chair, moving chair, taking our wallet, taking from wallet, mopping, and sweeping. Figure 12 shows the confusion matrix for the proposed framework when tested on the SYSU dataset. The overall accuracy for the classifier was 0.954. The Precision and Recall were 0.956 and 0.958, respectively. Table 5 reports on the precision, recall and the F-score of each class from the proposed framework. Table 6 indicates a comparison of HAR techniques indicating that the proposed framework outperforms baseline methods – according to accuracy – by 8.5%.

Confusion Matrix															
Output Class	Calling Phone	1073 10.2%	3 0.0%	0 0.0%	0 0.0%	0 0.0%	98 0.9%	0 0.0%	0 0.0%	0 0.0%	2 0.1%	6 0.1%	90.2% 9.8%		
	Drinking	4 0.0%	778 7.4%	0 0.0%	0 0.0%	0 0.0%	1 0.0%	3 0.0%	0 0.0%	0 0.0%	4 0.0%	2 0.0%	3 0.0%	97.9% 2.1%	
	Mopping	0 0.0%	0 0.0%	664 6.3%	1 0.0%	0 0.0%	0 0.0%	0 0.0%	1 0.1%	11 0.0%	0 0.0%	0 0.0%	0 0.0%	98.1% 1.9%	
	Moving Chair	0 0.0%	0 0.0%	0 0.0%	611 5.8%	3 0.0%	0 0.0%	0 0.0%	1 0.0%	0 0.0%	0 0.0%	0 0.0%	3 0.0%	98.9% 1.1%	
	Packing Backpack	0 0.0%	0 0.0%	0 0.0%	2 0.0%	1171 11.2%	0 0.0%	0 0.0%	0 0.0%	0 0.0%	0 0.0%	0 0.0%	3 0.0%	99.6% 0.4%	
	Playing Phone	11 0.1%	0 0.0%	0 0.0%	0 0.0%	0 0.0%	1131 10.8%	0 0.0%	0 0.0%	1 0.0%	6 0.1%	4 0.0%	0 0.0%	98.1% 1.9%	
	Pouring	1 0.0%	0 0.0%	0 0.0%	1 0.0%	1 0.0%	3 0.0%	798 7.6%	0 0.0%	0 0.0%	1 0.0%	0 0.0%	2 0.0%	98.9% 1.1%	
	Sitting Chair	2 0.0%	0 0.0%	0 0.0%	0 0.0%	3 0.0%	1 0.0%	4 0.0%	461 4.4%	0 0.0%	0 0.0%	0 0.0%	14 0.1%	95.1% 4.9%	
	Sweeping	0 0.0%	0 0.0%	23 0.2%	1 0.0%	0 0.0%	1 0.0%	0 0.0%	0 0.0%	648 6.2%	0 0.0%	0 0.0%	0 0.0%	98.3% 3.7%	
	Taking From Wallet	7 0.1%	1 0.0%	0 0.0%	0 0.0%	0 0.0%	1 0.0%	13 0.1%	0 0.0%	1 0.0%	0 0.0%	1203 11.5%	126 1.2%	4 0.0%	88.7% 11.3%
	Taking Our Wallet	2 0.0%	1 0.0%	0 0.0%	0 0.0%	1 0.0%	0 0.0%	7 0.0%	0 0.0%	4 0.0%	0 0.4%	39 5.4%	570 5.4%	13 0.1%	89.5% 10.5%
	Wearing Backpack	0 0.0%	0 0.0%	0 0.0%	0 0.0%	4 0.0%	3 0.0%	3 0.0%	0 0.0%	1 0.0%	0 0.0%	0 0.0%	2 0.0%	896 8.6%	98.6% 1.4%
		97.5% 2.5%	99.4% 0.6%	96.7% 3.3%	97.9% 2.1%	99.1% 0.9%	99.1% 0.9%	100% 0.0%	98.3% 1.7%	98.2% 1.8%	95.9% 4.1%	90.3% 9.7%	94.7% 5.3%	95.5% 4.5%	
Target Class															
	Calling Phone	Drinking	Mopping	Moving Chair	Packing Backpack	Playing Phone	Pouring	Sitting Chair	Sweeping	Taking From Wallet	Taking Our Wallet	Wearing Backpack			

Figure 12: Confusion Matrix for the proposed framework trained and tested on the SYSU-3D dataset.

Since the skeletal data provided in the SYSU-3D dataset can be regarded as noisy with large variations in viewpoint, the proposed framework brings significant performance improvement while showing how an activity recognition process can be insensitive to viewpoint variations.

TABLE 2: PROPOSED FRAMEWORK PERFORMANCE TESTED ON A SUBSET OF THE LBOROLDN HAR DATASET

	CARRYING BOXES	LYING DOWN	PUSHING A BOARD	RUNNING	SITTING ON A CHAIR	SITTING ON A STOOL	STANDING WHILE TEXTING	WALKING	WALKING WHILE TEXTING
PRECISION	1	1	0.988	0.726	0.881	0.968	0.9134	0.722	0.864
RECALL	0.959	1	1	0.688	0.890	0.957	0.989	0.776	0.795
F-SCORE	0.979	1	0.994	0.706	0.885	0.962	0.95	0.748	0.828
SAMPLE NO	74	103	86	73	101	94	104	72	81

TABLE 3: POINTNET PERFORMANCE TESTED ON A SUBSET OF THE LBOROLDN HAR DATASET

	CARRYING BOXES	LYING DOWN	PUSHING A BOARD	RUNNING	SITTING ON A CHAIR	SITTING ON A STOOL	STANDING WHILE TEXTING	WALKING	WALKING WHILE TEXTING
PRECISION	0.013	0.057	0.87	0	0.059	0	0	0	0
RECALL	0.019	0.714	0.11	0	0.038	0	NAN	NAN	NAN
F-SCORE	0.016	0.106	0.20	0	0.046	0	0	0	0
SAMPLE NO	74	103	86	73	101	94	104	72	81

TABLE 4: DNN PERFORMANCE TESTED ON A SUBSET OF THE LBOROLDN HAR DATASET

	CARRYING BOXES	LYING DOWN	PUSHING A BOARD	RUNNING	SITTING ON A CHAIR	SITTING ON A STOOL	STANDING WHILE TEXTING	WALKING	WALKING WHILE TEXTING
PRECISION	0.283	1	0.895	0.876	1	1	1	0.569	0.654
RECALL	0.677	1	0.987	0.771	0.980	1	0.990	0.577	0.441
F-SCORE	0.400	1	0.939	0.820	0.990	1	0.995	0.573	0.527
SAMPLE NO	74	103	86	73	101	94	104	72	81

TABLE 5: PROPOSED FRAMEWORK PERFORMANCE TESTED ON A SUBSET OF THE SYSU-3D DATASET

	CALLING PHONE	DRINKING	MOPPING	MOVING CHAIR	PACKING BACKPACK	PLAYING PHONE	POURING	SITTING CHAIR	SWEEPING	TAKING FROM WALLET	TAKING OUR WALLET	WEARING BACKPACK
PRECISION	0.975	0.993	0.966	0.979	0.990	0.896	1	0.982	0.981	0.958	0.802	0.947
RECALL	0.901	0.978	0.980	0.988	0.995	0.980	0.988	0.950	0.962	0.887	0.894	0.985
F-SCORE	0.937	0.986	0.973	0.983	0.993	0.936	0.994	0.966	0.972	0.921	0.846	0.966
SAMPLE NO	1100	783	687	624	1182	1262	798	469	660	1255	710	946

F. Visual Results of Proposed Framework

The purpose of the proposed algorithm is to assist an AV or AAL with HAR. Due to the multimodal requirements of the proposed framework, the results displayed here were a subset taken from the LboroLdnHAR dataset [3] while the experimental platform was stationary. All samples were gathered from indoor scenarios under controlled lighting.

The results presented in Figure 13 Scenario 1 (c), Scenario 2 (c), Scenario 3 (c) and Scenario 4 (c) indicate that the proposed HAR framework performed as desired. In all cases, the subject performing the activity was detected, segmented, and classified correctly.

For example, in Figure 13 Scenario 1 (a), the image shows a subject carrying a box. The first part of the proposed framework – object detector – identifies, with a high degree of confidence (99%), the location and region that the person occupies.

Figure 13 Scenario 1 (b) shows the 3D ROI identified by the object detector in the 2D image. Although there are unwanted artefacts located in the 3D ROI, the activity being performed by the subject gets correctly classified, as shown in Figure 13 Scenario 1 (c).

In Figure 13 Scenario 1 (a), there is a person – other than the subject – in the background. In this case, the object detector did not function as desired. This can be attributed to the design of the Faster-RCNN network, where input images with an aspect ratio of 360×360 are reduced down in size before processed.

It is important to remember, when processing images, the Faster-RCNN extracts features within the first few layers. When dealing with a small object in small images, the feature can, in effect, disappear in the middle of the network before detection occurs. In this case, the person in the background is not detected and therefore, their activity is never classified in the latter part of the network. Regardless of the framework missing this subject, the results show how multimodal machine learning can be successfully applied to HAR when using RGB data captured by a camera and point cloud data captured using Kinect® or LiDAR.

TABLE 6: COMPARISON OF ACTION RECOGNITION ACCURACY (%) WITH STATE-OF-THE-ART METHODS ON THE SYSU-3D DATASET

METHOD	YEAR	ACCURACY
ST-LSTM [72]	2018	73.4%
DYNAMIC SKELETONS [73]	2017	75.5%
DPRL [74]	2018	76.9%
VA-LSTM [75]	2017	77.5%
COMPLETE GR-GCN [76]	2019	77.9%
LOCAL+LGN [77]	2020	83.1%
ELEATT-GRU [78]	2020	85.7%
VA-FUSION [75]	2018	86.7%
SGNET [79]	2020	86.9%
PROPOSED FRAMEWORK	2020	95.4%

A similar but somewhat different issue occurs in Figure 13 Scenario 2 (a). In this image, the subject is standing while texting. The object detector correctly identifies the ROI for the subject but misses the person behind the computer due to occlusion. Contrary to this, the subject of interest was detected correctly, and the corresponding 3D ROI was extracted from the 3D data. Of particular interest is Figure 13 Scenario 2 (b). In this image, the subject is standing upright with hands out front, clasping a phone. This ROI is passed onto the classifier component of the proposed network before the correct result is overlaid onto the image shown in Figure 13 Scenario 2 (c).

It is important to note that the phone in Figure 13 Scenario 2 (b) occupies 19 3D points. Conversely, the person sitting in the background in Figure 13 Scenario 1 (a) occupies 79 pixels. Given that the object detector failed to identify the person in the background, it infers that the point cloud classifier is adept at

identifying point cloud features regardless of proximity.

A Further example of a dynamic activity is presented in Figure 13 Scenario 3 (a), where a subject is walking while texting. In this image, the subject was identified with 98% confidence. The corresponding 3D ROI was segmented using the ROI identified by the object detector is shown in Figure 13 Scenario 3 (b). As in Figure 13 Scenario 2 (b), the subject in this image had their hands out front clasping a phone. This region of interest was passed onto the classifier component before the correct result is overlaid onto the image shown in Figure 13 Scenario 3 (c).

Figure 13 Scenario 4 (a) shows a subject lying down. In this image, the area occupied by the subject, and some of the couch, was identified with 97% confidence. Figure 13 Scenario 4 (b) shows the 3D ROI identified by the object detector. In this image, the ROI contains the 3D data of the subject and a portion of the couch. When passed onto the classifier, the correct result was overlaid onto the image. Much like Figure 13 Scenario 1, where the subject was carrying a box, the object the subject is performing the activity with is incorporated into the classification processes.

V. DISCUSSION

In Figure 13, we presented four scenarios displaying the different section of the proposed framework for HAR. In each case, the model outputs remarkably accurate results. Furthermore, we found that the proposed framework can make an accurate prediction from partial data with few points.

For example, in Figure 13 Scenario 4 (b), the subject is in the supine position on a couch. The lack of point cloud data cannot be attributed to the instrument but the instrument principle of operation. Note the horizontal plains contacted by a laser – the further away, the greater the gap between plains.

In this case, points on the subject contacted by the laser are reduced to a single plain, making it difficult to classify. Moreover, given the obscurity and the absence of data points, humans trying to annotate the same data would find this a difficult task.

Figure 13: Scenario 1 showing the subject carrying a box. Scenario 2 showing the subject standing while texting. Scenario 3 showing the subject walking while texting. Scenario 4 showing the subject lying down. Images in the first column indicate the output of the object detector. The centre column shows the segmented point cloud data. The final column shows the output of the network and the associated class identified for the activity being performed.

Notwithstanding the successes, we do observe a few intrinsic issues to heed. The primary and most pronounced issue is to do with the frame rate of the LiDAR (5 HZ). While possible to increase the frame rate, reducing the number of points captured by a single scan. Although the proposed framework shows resilience to frugal data captures, there is a point where the network will miss classify. Although possible to reduce the frame rate below 5Hz and thus increase the data captured in a single frame. Doing so results in temporal ghosting of subjects moving at speed. When ghosting occurs, the RGB, point cloud and RGB-D data fall out of alignment. However, this can be managed with careful selection of the instruments frame rate.

We encountered other issues where the object detector failed to classify subjects at a distance from the camera. As commented in the results section, Figure 13 Scenario 1 (a) showed a subject performing the task and a person sitting on the couch. In this scenario, the object detector misses the person on the couch but identified the subject of interest. This issue was attributed to the distance the subject was from the RGB sensor. It is believed that this issue can be overcome by using sensors with a high-resolution image patch for far-away objects.

In terms of the point cloud classifier, it is important to note that the primary objective of all supervised learning framework is to discover the pattern linking the inputs X to the outputs Y , given a dataset D of size N [59], [60]. In the simplest terms, each input value has a dimension vector that represents the data we want the ML algorithm to learn. Commonly referred to as features, this representation is the understanding of X the algorithm develops during training. From the perspective of Y , the output can be anything, but the assumption is that it matches a categorical or nominal variable from the dataset used during training.

Generally, when using supervised learning methods for processing context data, a CNN obtains superior results with respect to traditional feature learning algorithms. Of course, in scenarios where data is sparse, a CNN can be outperformed. For example, when a CNN is trained on one dataset and expected to classify unfamiliar data that it has not been trained for, it will fail. However, when a CNN has been trained to recognise a specific action, such as falling, the framework should outperform all other forms of feature learning algorithms. Unfortunately, it is not always possible to acquire all the necessary data needed to train CNN to identify all specific actions. To that end, when we were collecting data, we focused on indoor activities where humans had a limited attention span. Had we included additional activities in the dataset used to train the proposed framework, it is clear from the results we presented in Section IV that they would have been classified with a high degree of accuracy.

In terms of comparison to other point cloud classifiers and their suitability for 3D feature learning, we looked at [62], [80], [81]. Researchers in [81] converted the point cloud data to a voxel grid array. How well the CNN acts on a voxel are dependent on the size of the grid – coarse or a fine. If applied to HAR, choosing a coarse grid leads to quantisation and a substantial loss of information, and a fine grid increases the computational cost.

PointNet was another approach to classification that directly consumes unordered point clouds [62], [80]. This context refers to the fact that the model does not assume any spatial relationships between features. This is not the case for HAR, which assumes neighbouring relations between the location of different points. Consequently, PointNet performs poorly when classifying HAR, regardless of how well it performs on classifying objects.

It should be noted that although the primary component of the proposed framework is a point cloud classifier. However, it is not the principal objective. The principal objective is to classify human activities robustly. With that in mind, we compared the proposed framework to a DNN framework for HAR [3], [46]. Although the DNN framework performed well, there are some intrinsic shortcomings. Primarily the issues with the DNN framework are to do with the type of data processed. Sensors that capture RGB-D data are notoriously susceptible to light changes [82] or interference from other sensors based upon structured light [83]. Consequently, the quality of light in the vision system environment is a key factor to its success. This is especially the case for networks that use sensors like the Kinect®. Similarly, networks that process RGB data are affected to a lesser degree D .

VI. CONCLUSION

This paper proposed a method of human activity recognition using multimodal machine learning to process RGB, RGB-D, and LiDAR data. The algorithm is composed of an object detector and a point cloud classification network. The object detector utilises a Faster R-CNN to identify a region of interest containing a person performing an activity. Translation and alignment of data types allow for the efficient segmentation of the point cloud data. The corresponding 3D region of interest is converted to a modified Fisher vector representation before a CNN classifies the activity. Performing Human Activity Recognition this way removes the need for wearable sensors to work in both indoor and outdoor environments. The technical challenges of the proposed framework lay with accurate alignment and translation of the different sensor modalities. Further technical challenges come in the form of sensor frequency of operation – which plays an integral role in perceiving the correct presence of subjects.

The object detection portion of the proposed framework accurately identifies the presence of subjects in an image, with an average precision of 95% and a log average miss rate of 0.1, respectively. The classification portion of the proposed framework accurately classified the activities performed by the subject with 90.3%. The precision and recall were 89.5% and 89.4%, respectively. Compared to an alternative point cloud classifier, the proposed framework significantly outperformed the next suitable alternative, PointNet, by a staggering 80.5%. Furthermore, it was found that the proposed framework outperformed a state-of-the-art HAR framework using a DNN to process RGB-D by 6.9%.

To benchmark the proposed framework, we conducted a systematic evaluation using the SYSU-3D dataset. The proposed method in this paper outperformed all reported

baseline methods – according to accuracy – by 8.5%. Since the skeletal data provided in the SYSU-3D dataset can be regarded as noisy with large variations in viewpoint, the proposed framework brings significant performance improvement. It should be noted that the SYSU-3D dataset is substantially larger than the LboraLdnHAR dataset by approximately 100 thousand samples. This sample size difference can account for the varying performance of the proposed framework when trained on different datasets

The implication of this research is machines that can better recognise human activities, without the need for wearable sensors, in varying environmental conditions and from various viewpoints. Indirectly this research will allow machines to recognise when a person needs help, autonomous vehicles to derive data-driven decision-making policies, and holds the potential to advance human-computer interaction to a new level of understanding.

REFERENCES

- [1] T. Bagautdinov, A. Alahi, F. Fleuret, P. Fua, and S. Savarese, “Social scene understanding: End-to-end multi-person action localization and collective activity recognition,” in *Proceedings - 30th IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017*, 2017.
- [2] X. Liu and C. S. Chua, “Multi-agent activity recognition using observation decomposed hidden Markov models,” *Image Vis. Comput.*, 2006.
- [3] M. Moencks, V. De-Silva, J. Roche, and A. Kondo, “Adaptive Feature Processing for Robust Human Activity Recognition on a Novel Multimodal Dataset,” *arXiv:1901.02858*, p. 13, 2019.
- [4] R. Q. Mínguez, I. P. Alonso, D. Fernández-Llorca, and M. Á. Sotelo, “Pedestrian Path, Pose, and Intention Prediction Through Gaussian Process Dynamical Models and Pedestrian Activity Recognition,” *IEEE Trans. Intell. Transp. Syst.*, vol. 20, no. 5, pp. 1803–1814, 2019.
- [5] De Silva V. Roche J. Kondo, A. V. De-Silva, J. Roche, and A. Kondo, “Robust Fusion of LiDAR and Wide-Angle Camera Data for Autonomous Mobile Robots,” *Sensors (Switzerland)*, vol. 18(8), no. 2730, 2018.
- [6] S. Ren, K. He, R. Girshick, and J. Sun, “Faster r-cnn: Towards real-time object detection with region proposal networks,” *IEEE Trans. Pattern Anal. Mach. Intell.*, 2017.
- [7] Y. Ben-Shabat, M. Lindenbaum, and A. Fischer, “3DmFV: Three-dimensional point cloud classification in real-time using convolutional neural networks,” *IEEE Robot. Autom. Lett.*, 2018.
- [8] B. P. Yuhas, M. H. Goldstein, and T. J. Sejnowski, “Integration of Acoustic and Visual Speech Signals Using Neural Networks,” *IEEE Commun. Mag.*, 1989.
- [9] X. Zhang *et al.*, “Emotion Recognition From Multimodal Physiological Signals Using a Regularized Deep Fusion of Kernel Machine,” *IEEE Trans. Cybern.*, pp. 1–14, 2020.
- [10] A. Hong *et al.*, “A Multimodal Emotional Human-Robot Interaction Architecture for Social Robots Engaged in Bidirectional Communication,” *IEEE Trans. Cybern.*, 2020.
- [11] Y. Yuan, Y. Lu, and Q. Wang, “Tracking as a Whole: Multi-Target Tracking by Modeling Group Behavior with Sequential Detection,” *IEEE Trans. Intell. Transp. Syst.*, 2017.
- [12] Y. Yuan, D. Wang, and Q. Wang, “Anomaly Detection in Traffic Scenes via Spatial-Aware Motion Reconstruction,” *IEEE Trans. Intell. Transp. Syst.*, 2017.
- [13] L. P. Morency and T. Baltrušaitis, “Multimodal machine learning: Integrating language, vision and speech,” in *ACL 2017 - 55th Annual Meeting of the Association for Computational Linguistics, Proceedings of the Conference Tutorial Abstracts*, 2017.
- [14] T. Baltrušaitis, C. Ahuja, and L. P. Morency, “Multimodal Machine Learning: A Survey and Taxonomy,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2019.
- [15] F. P. Such, V. Madhavan, E. Conti, J. Lehman, K. O. Stanley, and J. Clune, “Deep Neuroevolution: Genetic Algorithms Are a Competitive Alternative for Training Deep Neural Networks for Reinforcement Learning,” *arXiv:1712.06567*, 2017.
- [16] D. V. S. X. De Silva, W. A. C. Fernando, and S. T. Worrall, “Intra mode selection method for depth maps of 3D video based on rendering distortion modeling,” *IEEE Trans. Consum. Electron.*, 2010.
- [17] H. M. Le, C. Peter, and Y. Yue, “Data-Driven Ghosting using Deep Imputation Learning,” *MIT Sloan Sport. Anal. Conf.*, 2017.
- [18] V. Ghasemi and A. A. Pouyan, “Human activity recognition in ambient assisted living environments using a convex optimization problem,” in *Proceedings - 2016 2nd International Conference of Signal Processing and Intelligent Systems, ICSPIS 2016*, 2017.
- [19] H. Tabatabaee Malazi and M. Davari, “Combining emerging patterns with random forest for complex activity recognition in smart homes,” *Appl. Intell.*, 2018.
- [20] O. Banos, M. Damas, H. Pomares, F. Rojas, B. Delgado-Marquez, and O. Valenzuela, “Human activity recognition based on a sensor weighting hierarchical classifier,” *Soft Comput.*, 2013.
- [21] K. Gadzicki, R. Khamsehashari, and C. Zetzsche, “Early vs late fusion in multimodal convolutional neural networks,” in *Proceedings of 2020 23rd International Conference on Information Fusion, FUSION 2020*, 2020.
- [22] G. Chetty, M. White, M. Singh, and A. Mishra, “Multimodal activity recognition based on automatic feature discovery,” in *2014 International Conference on Computing for Sustainable Global Development, INDIACom 2014*, 2014.
- [23] F. Diederichs, T. Schuttke, and D. Spath, “Driver Intention Algorithm for Pedestrian Protection and Automated Emergency Braking Systems,” in *IEEE Conference on Intelligent Transportation Systems, Proceedings, ITSC*, 2015.
- [24] E. Rehder, H. Kloeden, and C. Stiller, “Head detection and orientation estimation for pedestrian safety,” in *2014 17th IEEE International Conference on Intelligent Transportation Systems, ITSC 2014*, 2014.
- [25] P. Gupta and T. Dallas, “Feature selection and activity recognition system using a single triaxial accelerometer,” *IEEE Trans. Biomed. Eng.*, 2014.
- [26] E. Garcia-Ceja, R. F. Brena, J. C. Carrasco-Jimenez, and L. Garrido, “Long-term activity recognition from wristwatch accelerometer data,” *Sensors (Switzerland)*, 2014.
- [27] Z. Chen, Q. Zhu, Y. C. Soh, and L. Zhang, “Robust Human Activity Recognition Using Smartphone Sensors via CT-PCA and Online SVM,” *IEEE Trans. Ind. Informatics*, 2017.
- [28] Z. Zhang and S. Poslad, “Improved use of foot force sensors and mobile phone GPS for mobility activity recognition,” *IEEE Sens. J.*, 2014.
- [29] Y. Kwon, K. Kang, and C. Bae, “Unsupervised learning for human activity recognition using smartphone sensors,” *Expert Syst. Appl.*, 2014.
- [30] M. Janidarmian, A. R. Fekr, K. Radecka, and Z. Zilic, “A comprehensive analysis on wearable acceleration sensors in human activity recognition,” *Sensors (Switzerland)*, 2017.
- [31] R. Igual, C. Medrano, and I. Plaza, “A comparison of public datasets for acceleration-based fall detection,” *Med. Eng. Phys.*, 2015.
- [32] A. Stisen *et al.*, “Smart devices are different: Assessing and mitigating mobile sensing heterogeneities for activity recognition,” in *SenSys 2015 - Proceedings of the 13th ACM Conference on Embedded Networked Sensor Systems*, 2015.
- [33] F. Attal, S. Mohammed, M. Dedabrishvili, F. Chamroukhi, L. Oukhellou, and Y. Amirat, “Physical human activity recognition using wearable sensors,” *Sensors (Switzerland)*, 2015.
- [34] D. M. Karantonis, M. R. Narayanan, M. Mathie, N. H. Lovell, and B. G. Celler, “Implementation of a real-time human movement classifier using a triaxial accelerometer for ambulatory monitoring,” *IEEE Trans. Inf. Technol. Biomed.*, 2006.
- [35] L. Atallah, B. Lo, R. King, and G. Z. Yang, “Sensor placement for activity detection using wearable accelerometers,” in *2010 International Conference on Body Sensor Networks, BSN 2010*, 2010.
- [36] J. Y. Yang, J. S. Wang, and Y. P. Chen, “Using acceleration measurements for activity recognition: An effective learning algorithm for constructing neural classifiers,” *Pattern Recognit. Lett.*, 2008.
- [37] M. M. Hamdi, M. I. Awad, M. M. Abdelhameed, and F. A. Tolbah, “Lower limb gait activity recognition using Inertial Measurement Units for rehabilitation robotics,” in *Proceedings of the 17th International Conference on Advanced Robotics, ICAR 2015*, 2015.
- [38] X. Hu, C. Yao, and G. S. Soh, “Performance evaluation of lower limb ambulatory measurement using reduced Inertial Measurement Units and 3R gait model,” in *IEEE International Conference on Rehabilitation Robotics*, 2015.

- [39] S. I. Patel and M. J. Souter, "Equipment-related electrocardiographic artifacts: Causes, characteristics, consequences, and correction," *Anesthesiology*, 2008.
- [40] A. Vaidya, R. Bentley-Lewis, X. Jeunemaitre, G. K. Adler, and J. S. Williams, "Dietary sodium alters the prevalence of electrocardiogram determined left ventricular hypertrophy in hypertension," *Am. J. Hypertens.*, 2009.
- [41] G. M. Stern and A. S. D. Spiers, "Hypotension Caused by L-Dopa," *Br. Med. J.*, 1970.
- [42] S. U. Park, J. H. Park, M. A. Al-Masni, M. A. Al-Antari, M. Z. Uddin, and T. S. Kim, "A Depth Camera-based Human Activity Recognition via Deep Learning Recurrent Neural Network for Health and Social Care Services," in *Procedia Computer Science*, 2016.
- [43] K.-T. Song and W.-J. Chen, "Human activity recognition using a mobile camera," in *2011 8th International Conference on Ubiquitous Robots and Ambient Intelligence (URAI)*, 2011, pp. 3–8.
- [44] K. He, X. Zhang, S. Ren, and J. Sun, "ResNet," *Proc. IEEE Comput. Soc. Conf. Comput. Vis. Pattern Recognit.*, 2016.
- [45] S. Mukherjee, L. Anvitha, and T. M. Lahari, "Human activity recognition in RGB-D videos by dynamic images," *Multimed. Tools Appl.*, 2020.
- [46] E. Cippitelli, S. Gasparrini, E. Gambi, and S. Spinsante, "A Human Activity Recognition System Using Skeleton Data from RGBD Sensors," *Comput. Intell. Neurosci.*, 2016.
- [47] H. Feng, S. Wang, H. Xu, and S. S. Ge, "Object Activity Scene Description, Construction, and Recognition," *IEEE Trans. Cybern.*, pp. 1–11, 2019.
- [48] M. M. Arzani, M. Fathy, A. A. Azirani, and E. Adeli, "Switching Structured Prediction for Simple and Complex Human Activity Recognition," *IEEE Trans. Cybern.*, 2020.
- [49] J. Byun, K.-I. Na, M. Noh, and S. Kim, "ESTRO: Design and development of intelligent autonomous vehicle for shuttle service in the ETRI," in *Proceedings of the Workshop Planning Perception Navigation Intelligent Vehicles*, 2012.
- [50] N. Suganuma and T. Uozumi, "Development of an autonomous vehicle - System overview of test ride vehicle in the Tokyo Motor Show 2011," in *Proceedings of the SICE Annual Conference*, 2012.
- [51] A. Mancini, E. Frontoni, and P. Zingaretti, "Embedded Multisensor System for Safe Point-to-Point Navigation of Impaired Users," *IEEE Trans. Intell. Transp. Syst.*, 2015.
- [52] T. Pallejā, M. Tresanchez, M. Teixidó, and J. Palacin, "Bioinspired electronic white cane implementation based on a LIDAR, a tri-axial accelerometer and a tactile belt," *Sensors*, 2010.
- [53] B. Centre and A. R. Sensing, "LiDAR – Overview of Technology , Applications , Market Features & Industry," *Victoria*, 2006.
- [54] R. Saracco, "LIDAR hits mass market," *IEEE Future Directions*, 2020.
- [55] R. Saracco, "Light field Technology – the future of photography?," *IEEE Future Directions*, 2020.
- [56] F. Luo, S. Poslad, and E. Bodanese, "Temporal convolutional networks for multi-person activity recognition using a 2D LIDAR," *IEEE Internet Things J.*, 2020.
- [57] J. Wang, Y. Chen, S. Hao, X. Peng, and L. Hu, "Deep learning for sensor-based activity recognition: A survey," *Pattern Recognit. Lett.*, 2019.
- [58] H. F. Nweke, Y. W. Teh, M. A. Al-garadi, and U. R. Alo, "Deep learning algorithms for human activity recognition using mobile and wearable sensor networks: State of the art and research challenges," *Expert Systems with Applications*, 2018.
- [59] N. Pittaras, F. Markatopoulou, V. Mezaris, and I. Patras, "Volumetric and Multi-View CNNs for Object Classification on 3D Data Supplementary Material," *Lect. Notes Comput. Sci. (including Subser. Lect. Notes Artif. Intell. Lect. Notes Bioinformatics)*, 2017.
- [60] Z. Wang, W. Zhan, and M. Tomizuka, "Fusing Bird's Eye View LIDAR Point Cloud and Front View Camera Image for 3D Object Detection," in *IEEE Intelligent Vehicles Symposium, Proceedings*, 2018.
- [61] Y. Cai, T. Zhang, H. Wang, Y. Li, Q. Liu, and X. Chen, "3D Vehicle Detection Based on LiDAR and Camera Fusion," *Automot. Innov.*, 2019.
- [62] C. R. Qi, L. Yi, H. Su, and L. J. Guibas, "PointNet++: Deep hierarchical feature learning on point sets in a metric space," in *Advances in Neural Information Processing Systems*, 2017.
- [63] C. R. Qi, H. Su, K. Mo, and L. J. Guibas, "PointNet: Deep learning on point sets for 3D classification and segmentation," in *Proceedings - 30th IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017*, 2017.
- [64] A. Krizhevsky, I. Sutskever, and H. Geoffrey E., "Imagenet," *Adv. Neural Inf. Process. Syst.* 25, 2012.
- [65] D. W. James, J. Eckermann, B. Fawsi, and J. Sienz, "Point cloud data from Photogrammetry techniques to generate 3D Geometry," in *Conference of the Association for Computational Mechanics in Engineering*, 2015, p. 4.
- [66] R. A. Fisher, "Inverse Probability," *Math. Proc. Cambridge Philos. Soc.*, 1930.
- [67] G. Zeng, Y. He, Z. Yu, X. Yang, R. Yang, and L. Zhang, "InceptionNet/GoogLeNet - Going Deeper with Convolutions," *Cypr*, 2016.
- [68] C. Szegedy *et al.*, "Going deeper with convolutions," in *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 2015.
- [69] M. Everingham, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman, "The pascal visual object classes (VOC) challenge," *Int. J. Comput. Vis.*, 2010.
- [70] S. Ceri *et al.*, *An Introduction to Information Retrieval*. Cambridge University Press, 2013.
- [71] P. Bhatia, *Data Mining and Data Warehousing: Principles and Practical Techniques*. Cambridge University Press, 2019.
- [72] J. Liu, A. Shahroudy, D. Xu, A. C. Kot, and G. Wang, "Skeleton-based action recognition using spatio-temporal LSTM network with trust gates," *arXiv*, 2017.
- [73] J. F. Hu, W. S. Zheng, J. Lai, and J. Zhang, "Jointly Learning Heterogeneous Features for RGB-D Activity Recognition," *IEEE Trans. Pattern Anal. Mach. Intell.*, 2017.
- [74] Y. Tang, Y. Tian, J. Lu, P. Li, and J. Zhou, "Deep Progressive Reinforcement Learning for Skeleton-Based Action Recognition," in *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 2018.
- [75] P. Zhang, C. Lan, J. Xing, W. Zeng, J. Xue, and N. Zheng, "View Adaptive Recurrent Neural Networks for High Performance Human Action Recognition from Skeleton Data," in *Proceedings of the IEEE International Conference on Computer Vision*, 2017.
- [76] X. Gao, W. Hu, J. Tang, J. Liu, and Z. Guo, "Optimized skeleton-based action recognition via sparsified graph regression," in *MM 2019 - Proceedings of the 27th ACM International Conference on Multimedia*, 2019.
- [77] Q. Ke, M. Bennamoun, H. Rahmani, S. An, F. Sohel, and F. Boussaid, "Learning Latent Global Network for Skeleton-Based Action Prediction," *IEEE Trans. Image Process.*, 2020.
- [78] P. Zhang, J. Xue, C. Lan, W. Zeng, Z. Gao, and N. Zheng, "EleAtt-RNN: Adding attentiveness to neurons in recurrent neural networks," *IEEE Trans. Image Process.*, 2020.
- [79] P. Zhang, C. Lan, W. Zeng, J. Xing, J. Xue, and N. Zheng, "Semantics-guided neural networks for efficient skeleton-based human action recognition," in *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 2020.
- [80] C. R. Qi, W. Liu, C. Wu, H. Su, and L. J. Guibas, "Frustum PointNets for 3D Object Detection from RGB-D Data," in *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 2018.
- [81] Y. Zhou and O. Tuzel, "VoxelNet: End-to-End Learning for Point Cloud Based 3D Object Detection," in *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 2018.
- [82] J. Jancek, D. Aleinikava, and G. M. Mirsky, "Optimizing kinect® depth sensing using dynamic polarization," in *Proceedings of the Conference on Integrating Technology into Computer Science Education, ITiCSE*, 2017.
- [83] R. Martin Martin, M. Lorbach, and O. Brock, "Deterioration of depth measurements due to interference of multiple RGB-D sensors," in *IEEE International Conference on Intelligent Robots and Systems*, 2014.