

Received December 30, 2020, accepted January 16, 2021, date of publication January 20, 2021, date of current version February 1, 2021.

Digital Object Identifier 10.1109/ACCESS.2021.3053084

Machine-Generated Hierarchical Structure of Human Activities to Reveal How Machines Think

MAHSUN ALTIN¹, FURKAN GÜRSOY², AND LINA XU³

¹Department of Computer Engineering, Abdullah Gül University, 38080 Kayseri, Turkey

²Department of Management Information Systems, Boğaziçi University, 34342 Istanbul, Turkey

³Department of Computer Science, University College Dublin, Dublin 4, Ireland

Corresponding author: Furkan Gürsoy (furkan.gursoy@boun.edu.tr)

ABSTRACT Deep-learning based computer vision models have proved themselves to be ground-breaking approaches to human activity recognition (HAR). However, most existing works are dedicated to improve the prediction accuracy through either creating new model architectures, increasing model complexity, or refining model parameters by training on larger datasets. Here, we propose an alternative idea, differing from existing work, to increase model accuracy and also to shape model predictions to align with human understandings through automatically creating higher-level summarizing labels for similar groups of human activities. First, we argue the importance and feasibility of constructing a hierarchical labeling system for human activity recognition. Then, we utilize the predictions of a black box HAR model to identify similarities between different activities. Finally, we tailor hierarchical clustering methods to automatically generate hierarchical trees of activities and conduct experiments. In this system, the activity labels on the same level will have a designed magnitude of accuracy and reflect a specific amount of activity details. This strategy enables a trade-off between the extent of the details in the recognized activity and the user privacy by masking some sensitive predictions; and also provides possibilities for the use of formerly prohibited invasive models in privacy-concerned scenarios. Since the hierarchy is generated from the machine's perspective, the predictions at the upper levels provide better accuracy, which is especially useful when there are too detailed labels in the training set that are rather trivial to the final prediction goal. Moreover, the analysis of the structure of these trees can reveal the biases in the prediction model and guide future data collection strategies.

INDEX TERMS Hierarchical labeling, human activity recognition, machine learning, privacy preservation, video processing.


I. INTRODUCTION

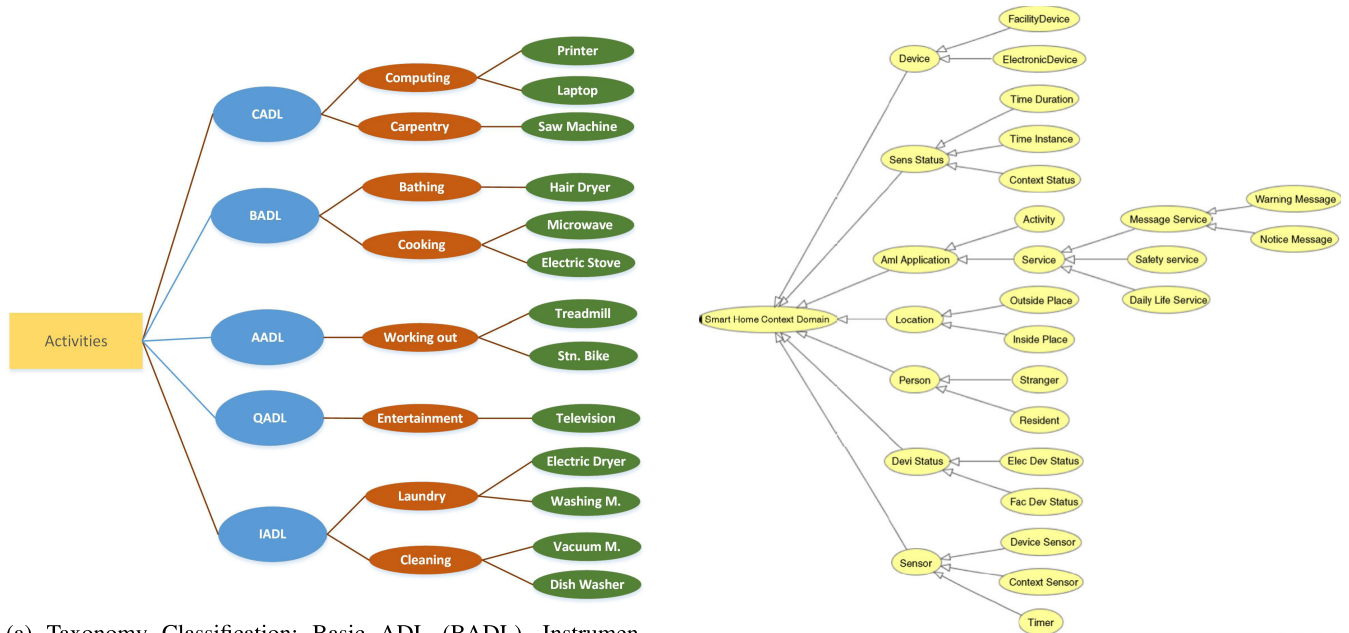
Video processing attracts a huge demand and interest from both academia and industry due to its ability to unlock the intelligence in many domains such as surveillance, gaming, autonomous vehicles, medical imaging, human activity recognition, and alike. Therefore, video processing and analysis have become a predominant research topic in the artificial intelligence (AI) and machine learning (ML) field. The performance of video processing tasks has been brought to a new level with the utilization of deep learning (DL).

In existing ML models for video processing and labeling approaches, we have problems such as (i) highly correlated labels and low accuracy for activities involving richer context

(e.g., labels with very specific definitions when the number of labels is high) (ii) too detailed prediction results that can violate privacy (iii) trivial labels that have no meaning for a specific use scenario.

First, existing ML models increasingly try to find specific labels for a larger number of activities. Their accuracy may become lower since it will be more challenging to distinguish between different but similar activities. Without understanding the correlation between the observed data, more data collection can be meaningless [1]. Besides, without understanding the fundamentals of the model (often referred to as the machine), it is hard to redesign the model structure. Given these issues, the explainable AI approach [2]–[4] aims to understand how the machine thinks while possibly using such understanding to improve prediction accuracy or usefulness in real-world scenarios with various constraints.

The associate editor coordinating the review of this manuscript and approving it for publication was Yiming Tang .



(a) Taxonomy Classification: Basic ADL (BADL), Instrumental ADL (IADL), Extended ADL (EADL), and Advanced ADL (AADL) (Adapted from [5]).

(b) Web Ontology Language (OWL) based Ontology for ADL Recognition in Smart Homes (Adapted from [6]).

FIGURE 1. Examples of existing approaches for ADL understanding and classification.

Most existing work has been focusing on training data set and tuning fine models, while limited attention has been given to the post-processing related to the model predictions results.

Second, privacy has become a highly concerning issue in many human-computer systems. In the context of monitoring, privacy has been a controversial and focal point, especially with the enforcement of certain laws such as the General Data Protection Regulation (GDPR). The massive video image collections and the intelligence of processing ability have already brought many scary thoughts to the public. For example, in the healthcare domain, being able to predict human’s daily activities of living (ADL) and perform auto logging will greatly help to build an evidence-based medicine system. However, ADL classification through video can be extremely invasive without information filter control. With existing ML models, it is hard to control and balance between the prediction accuracy and privacy. Hence, there should be a way to trade off between the details presented from the video analysis results and privacy while retaining the effectiveness of the system and satisfying the privacy expectations of people.

Third, with the increasing number of possible activity labels, some of those labels have no use in many real-world applications. For instance, the label *alligator wrestling* from *Kinetics-600* dataset will not be relevant in most use cases. Hence, such models and their results should be processed with an appropriate strategy that accounts for the trivial labels.

The traditional way of understanding structure and relationships among human activities is based on taxonomy

or ontology knowledge graphs, as visually exemplified in Figure 1. However, in existing approaches, the structure of these taxonomies (i.e., activity trees) are generated from the human’s perspective. Yet, if we want to apply ML models for analyzing videos to identify ADL and evaluate the models’ intelligence in order to redesign the model, it is also important to understand the classification logic from the machine’s perspective. This is essential to improve and tailor the ML models’ performance in a clear direction.

In order to address the above-stated problems, we present a hierarchical labeling system for human activity recognition. Such a system is inspired from the machine’s perspective and aims to provide adjustable accuracy and privacy for different use scenarios. The concept of the system is illustrated in Figure 2. Each level has a designed accuracy range and context richness. Based on the use case requirements on privacy, the system can be set at a specific level. For example, in an edge-cloud configuration, the data collection and analysis can be done on the edge and only summarized results can be transmitted to the cloud so that the private information is not seen or stored by the system operator. Since the hierarchical structure is generated through observations from the machine’s perspective, the prediction accuracy will increase as LEVEL moves from bottom to up. Moreover, the similarities between labels as understood by the machine will be revealed by the hierarchical tree, which in turn enables model understanding and may guide future data collection strategies.

To the best of knowledge, this is the first study to propose a fully machine-generated human activity hierarchy to

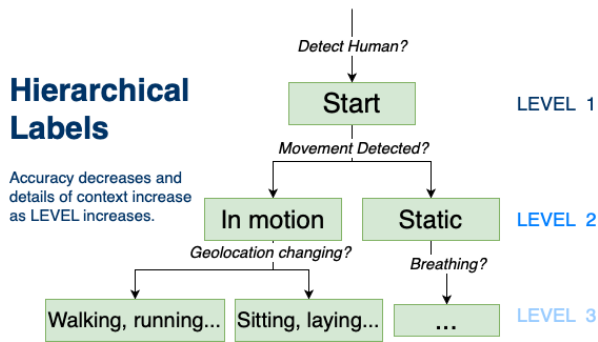


FIGURE 2. Analysis supporting a hierarchical labeling system.

overcome above-stated problems and reveal the biases of a HAR machine, and explore how it sees the world; eventually allowing and guiding researchers and practitioners to design appropriate mechanisms in real-world use scenarios and also to improve the model accuracy. In the next section, we provide a brief review of methods and datasets for human activity recognition, present our motivation, and discuss the feasibility of our approach. In Section 3, we describe our methods for automatically generating the hierarchical trees and demonstrate our strategy utilizing a chosen deep learning model and several public datasets. We present experimental results and provide quantitative and qualitative analyses in Section 4. Finally, we present our final remarks and a brief discussion on future directions.

II. BACKGROUND

A. RELATED WORK

Human activity recognition aims to analyze and detect different human actions from various data sources such as sensor data or multimedia data. It finds application areas in security and surveillance, healthcare, autonomous driving, robotics, smart home, entertainment, and others. Generally speaking, the literature on HAR is clustered around two topics: the more traditional sensor-based HAR and vision-based HAR. For vision-based HAR, earlier methods include more traditional machine learning algorithms such as decision trees or support vector machines which require hand-engineered features. Recently, deep learning-based models gained popularity since they do not require a time-consuming manual feature generation process and show superior performance. In particular, sequential learning methods (e.g., long-short term memory) and convolutional neural networks play a key role in the architecture of the state-of-the-art deep learning-based HAR models [7].

Providing a comprehensive review of the developments in and the current state of HAR research is not possible given the limited space and is beyond the scope of this work. We refer interested readers to the recent comprehensive survey paper by Dang et al. [8] that provides a relatively concise review of hundreds of relevant studies among which exist more than 10 other HAR survey papers. However, here, we briefly describe the model used in this study. Two-Stream Inflated




FIGURE 3. Two-stream inflated 3D ConvNet model, input and output.

3D Convnet (*I3D*) [9] is a state-of-the-art deep learning-based computer vision model used in human action recognition. As illustrated in Figure 3, taking an RGB video as the input, the model produces a series of probabilities corresponding to the labels. The label with the highest probability can be considered as the prediction result for the video. Alternatively, the top k results can be considered together as top predictions.

Many human activity datasets of different formats and coverage have been employed in the literature. It is beyond the scope of this work to list and present datasets. There are studies that specifically review and compare HAR datasets in the literature [10]–[12]. Instead, here, we only present a summary of several public, popular, and relevant video datasets in Table 1. *Kinetics* has been first published as *Kinetics-400* [9]. It has 400 different classes and 306, 245 videos all with a length of 10 seconds. *Kinetics-600* extended it to cover 600 classes 495, 547 videos [13]. Recently, *Kinetics-700* was made available [14]. All of the video labels in Kinetics dataset are homogeneous and no hierarchical labels are considered. Kuchne et al. [15] published *HMDB51* which covers 51 different human activities with 6, 766 videos, with a minimum 1 second long length. All the videos are also labeled with meta labels, which can be used to create higher classes among those 51 classes. Soomro et al. [16] published *UCF101* which covers 101 different classes and 13, 320 videos. It applied a simplified approach of ontology where the action categories can be divided into five types: human-object interaction, body-motion only, human-human interaction, playing musical instruments, and sports. *ActivityNet* was originally published as *ActivityNet-100* [17] and later updated as *ActivityNet-200*, containing 849 hours long videos in total. It has organized the labels into a relatively richer hierarchical tree through the expert taxonomy approach. Similarly, *FCVID* [18] categorizes activities in 11 high-level groups. Lastly, *COIN* dataset [19] employs a similar taxonomy classification with 3-stages (domain, task, and step) and contains 11, 827 instructional videos related to 180 different tasks in 12 domains.

As previously illustrated in Section I, existing studies focusing on human activity hierarchies employs expert generated hierarchies such as in [5], [6], [20]. Long et al. [21] utilizes action hierarchies represented in a hyperbolic space rather than simpler tree representations, where hierarchies are based on some external expert guidelines from *ActivityNet*. Gaidon et al. [22] represent individual videos as a hierarchy

TABLE 1. Open video datasets.

Name	# of Videos	# of Classes	Maintained By	Hierarchical Label
Kinetics-400.	306,245	400	Deepmind	No
Kinetics-600.	495,547	600	Deepmind	No
HMDB51	6,766	51	Serre-Lab	Yes, meta labels
UCF101	13,320	101	U. of Central Florida	Yes, simplified ontology
ActivityNet 200	27,801	203	ActivityNet	Yes, taxonomy
FCVID	91,223	239	Fudan University	Yes, taxonomy
COIN	11,827	180	Tsinghua University	Yes, taxonomy

of activities based on spatio-temporal trajectories but do not provide a global taxonomy of activities. Our strategy differs from these and brings a novel approach to understand the behavior and bias of an HAR machine and leverage the power of human intelligence to improve it.

B. MOTIVATION AND FEASIBILITY

The ultimate goal of AI is to train the machine to think in a consistent way with how a human thinks, particularly in tasks where humans are good at such as human activity recognition. Unfortunately, we are currently far from that and there is still a large gap between how the machines think and how we humans think. In order to reduce this gap, it is necessary to assess where the machine's intelligence is through understanding the analysis results from the machine's perspective. This awareness will be a key to enhance the learning model and it can also provide directions for future data collection. In order to train the machine and collect more data in an efficient and effective manner, one needs to understand the correlation between the existing videos and address the limitations and boundaries. Most existing models are mainly concerned with the accuracy for each label and video. We think that introducing a hierarchical prediction system can better serve the goal. Moreover, one huge concern currently about video processing is privacy invasion. Limited work concerned itself with the privacy problem when designing an analyzing/predicting machine learning model for ADL recognition.

In order to modify existing deep learning models to output prediction results through a hierarchical labeling structure, firstly we need to analyze the current prediction results and prove that there is such a possibility. This requires us to understand the human activities from the machine's perspective, rather than from our human intuition. The machine learns about the world through the dataset that is fed to them, more specifically, videos and the labels that are associated with them. The model might be biased and contradicting to our intuition due to the explicit or implicit limitations of the dataset. For example, from the perspective of a human or a real-world use case, we may think videos of *brushing teeth* and *flossing* are quite similar. However, an ML model may classify them quite differently and identify that *flossing* has a higher similarity to *putting makeup*. This may simply because the training dataset does not have enough number of diverse videos for these labels, hence, the model cannot identify this activity precisely.

We have examined several existing video analysis deep learning models and discover that the total number of labels are huge, ranging to more than 600 labels. Some of the labels are revealing great details of human activities. In many scenarios, too detailed labeling is not necessary and also invasive to privacy. For example, due to privacy issues, there would be different requirements for the analysis results and the level of details. Besides, too detailed labels make training those models become difficult since sourcing appropriately rich and diverse training data is challenging. Due to the limited amount of training data, the prediction accuracy for individual labels can be extremely low, which provides little meaning. Some of the labels have high correlations, which makes prediction even more difficult and the ML model can be easily confused. Some of the labels are mostly trivial and make little sense in reality, especially in daily activities monitoring. To summarize, there are several existing problems accosted with existing video processing deep learning models:

- Too many labels with only homogeneous organizing
- Low accuracy for top 1 prediction
- High correlation between many labels
- Too detailed labeling invading privacy

The Two-Stream I3D model trained on *ImageNet* [23] and *Kinetics* dataset yields the following performance on *Kinetics-400*. The accuracy for the Top 1 prediction is 74.2% whereas the accuracy for the sum of top 5 predictions can reach 91.3%. Some activities with more details/context involved are harder to predict than the ones with less detail. For example, *tango dancing* is harder to identify comparing with *jogging*. Furthermore, *tango dancing* is reflecting more details related to a scenario than *dancing*.

This motivates us to investigate into the similarity of the labels and build a hierarchical tree structure of them by ML techniques. The depth of the tree indicates the accuracy and level of details of the prediction result to fulfill a system's privacy and accuracy requirements. In order to construct such a hierarchical tree for ADL, we need to understand the correlations between the labels and this understanding needs to be from the machine's perspective. Through this way, we can discover which parts of the machine's thinking are contradicting to our human common sense. Therefore we can fill the gap via specifically directed future improvements. Then, for privacy concerns, we need to understand the connection between privacy and accuracy. We can choose a method to cluster activity labels into a hierarchical structure to enable

the high prediction accuracy at desired, privacy-preserving levels.

III. HIERARCHICAL LABELLING SOLUTION

The methodology in this study is designed to generate a hierarchy labeling tree for human activities and can be conducted in two consecutive steps. First, similarity measures are devised to quantify the resemblance between different activities through their co-occurrences in the predictions of a video classification model. This resemblance is measured and assessed from the machine's perspective rather than the human's. Then, considering the quantified similarities between activities as the distance metric, a set of various hierarchical clustering techniques are applied to automatically generate the activity label trees. Through analyzing the generated tree from the human's perspective, we can gain insights into the model and the dataset.

A. SIMILARITY SPACE

In order to construct the similarity dataset, we have followed the following procedures.

1) COMPOSED DATASET

A collective dataset with 15,938 videos are formed from public datasets including *UCF101*, *HMDB51*, and *Kinetics-600*. Herein we refer this dataset as the *ComposedDS*. The whole *UCF101* dataset which consists of 12,435 videos is taken. 3456 videos are chosen randomly from *HMDB51*. However, the selected videos from the two datasets did not have any videos that resembled the 6 of the 600 activities from *Kinetics*, which hints at the discrepancy between different datasets. To ensure that all labels show up at least once in the top-5 predictions of the *I3D* model, 47 videos of these 6 activities are taken from the test set of *Kinetics-600*.

2) VIDEO CLASSIFIER

The *I3D* model pre-trained on the training set of *Kinetics-600* is used to make predictions for each of video from *ComposedDS*. The model has an output vector with 600 values through a softmax function. Each value in the vector represents the predicted probability of the corresponding label. The prediction results for all the videos from *ComposedDS* are recorded and used to form a generated dataset with a size of [15938, 600], herein named *GeneratedDS600*. It is important to note that the *I3D* model was not trained on any of the videos from *ComposedDS* and the original ground-truth labels of *ComposedDS* differs from that of *Kinetics-600*. Therefore, the dataset that the *I3D* model is trained on does not overlap with the dataset that is used to generate *GeneratedDS*. For each video, the top five predictions of *ComposedDS* are recorded as an individual set, named *GeneratedDS5* with a size of [15938, 5]. For an example video, the top five predictions are stored in *GeneratedDS5* as shown in Table 2, originally followed by No. 6 *passing American football (AF)*, *in game* with a probability of 0.02.

TABLE 2. *I3D* predictions for a sample video.

Predicting Label	Predicting Score (Sorted)
shooting goal (soccer)	0.58
passing soccer ball	0.18
playing field hockey	0.06
kicking soccer ball	0.06
juggling soccer ball	0.04

A really interesting fact we have observed from Table 2 and many other instances from *GeneratedDS5*, is that most of the time, the activities that the machine thinks are correlated, are also correlated from our human's understanding. This gives us a good start. *I3D* model is recognizing the world similarly as a human to a certain degree. However, it is not always the case. As in Table 2, the machine thinks that *playing field hockey* is similar to other soccer activities and ranks better than few other soccer activities in this case. This mistake is not dramatically disappointing since the probabilities associated with those labels are not high. With a simple analogy, if a child has made this mistake, blindly forcing them to remember this activity is not an efficient way of education. Finding the correlation between the activities and teaching the child to learn the similarities and the distinctions is a better approach. Therefore, teaching the child to distinguish a soccer field from the hockey field may avoid this mistake in the future. The same is true when training machines.

3) SIMILARITY CALCULATION

To teach the machine to be more intelligent, we want to adopt a more systematic way rather than just randomly collecting more data and perform more training. We need to understand why the machine made this mistake. Like teaching a child, firstly, we need to find out, which labels that the machine thinks are highly similar. We propose a similarity measure based on the element co-occurrences in *GeneratedDS5*. We propose to use two different measures to quantify the extent of co-occurrence: *confidence* [24] (Eq. 1) and *lift* [25] (Eq. 2).

$$C_{ij} = P(i|j) = \frac{P(\{i, j\})}{P(j)} \quad (1)$$

$$L_{ij} = \frac{P(i|j)}{P(i)} = \frac{P(\{i, j\})}{P(i)P(j)} \quad (2)$$

Confidence C_{ij} measures the probability of i being in the set given that j already exists. C_{ij} can range from 0 to 1 and a higher value indicates a greater similarity between the activity i and j . It is an asymmetric measure; thus, $C_{ij} \neq C_{ji}$ usually. Moreover, if i is a popular activity in the dataset, it tends to co-occur more with others, resulting in having higher *confidence* values.

Lift L_{ij} corrects C_{ij} for the overall popularity of i over all sets. When it equals 1, it indicates the complete independence of events i and j , i.e., the extent of their co-occurrences is by random chance. Like *confidence*, the higher values indicates greater similarity. Unlike the former which is limited to the

[0, 1] range, it can take any non-negative value and it is a symmetric metric; thus, $L_{ij} = L_{ji}$.

We can consider that high *confidence* or *lift* value indicate high similarity between two activities from the machine's perspective. However, we are also interested in the similarity structure among all activities. We use *confidence* and *lift* as the distance metrics to cluster activities into subgroups. The information for the subgroups is important for testing a machine and further training. Using the same analogy, when asked to describe an activity, a child might say *playing soccer*. This may be a wrong answer, but you can still get the general concept that people are playing a certain sport. If people are not playing soccer but are playing sports, then you can still trust the explanation of the child at a certain level. If people are not even playing sports, then you cannot rely on the responses of the child at all. This hints that there are more fundamental issues about how the child is taught to distinguish different activities. Likewise, in the machine intelligence case, we want to know what is the probability associated with playing sports and what is the probability associated with playing soccer. Therefore, hierarchical clustering for activities is a good strategy to assess the intelligence of a machine and also can help with teaching a machine to learn.

For hierarchical clustering, the distance between i and j must be symmetric, i.e., $D_{ij} = D_{ji}$. *Lift* is symmetric whereas *confidence* is not. We define the distance D_{ij} between two activities i and j in two alternative ways based on *confidence* and *lift* as shown in Equation 3 and 4 respectively and ensure that both are generalized to the range [0, 1] for consistency.

$$D_{ij}^C = 1 - \sqrt{C_{ij}C_{ji}} \quad (3)$$

$$D_{ij}^L = 1 - \frac{L_{ij} - \min(L)}{\max(L) - \min(L)} \quad (4)$$

In obtaining a single symmetric measure from two asymmetric *confidence* measures, geometric mean is chosen over arithmetic mean on purpose to award reciprocal similarities where activity i co-occurs with activity j , and j also co-occurs with i . The measure obtained by geometric mean is referred to as the *cosine* measure whereas the measure that can be obtained by the arithmetic mean is referred to as *Kulczynski* measure. *Cosine* measure is *null-invariant*, i.e., it is not affected by the sets that do not contain any of the examined elements (i.e., null sets). On the other hand, *lift* is sensitive to the number of such null sets, which [26, Chapter 6] provides a comparative discussion on different measures in the context of frequent itemset mining. As a demonstrating example, Table 3 lists each activity that is a member in at least 5% of all sets where *playing basketball* is also a member, along with the values calculated by the similarity measures employed in this study.

B. HIERARCHICAL CLUSTERING

Clustering is the process of creating multiple groups (i.e., clusters) from the given objects such that the objects are similar to those in the same clusters and dissimilar to those

TABLE 3. Top items co-occurring with j : *playing basketball*.

i	$P(i j)$	$P(j i)$	C_{ij}	L_{ij}
shooting basketball	0.76	0.87	0.81	23.5
dunking basketball	0.64	0.87	0.74	23.5
dribbling basketball	0.54	0.82	0.67	22.1
playing netball	0.33	0.58	0.44	15.5
playing volleyball	0.20	0.63	0.36	17.0
juggling soccer ball	0.19	0.30	0.24	8.0
passing AF, not in game	0.14	0.51	0.27	13.8
throwing ball (not baseball or AF)	0.11	0.25	0.17	6.7
dodgeball	0.08	0.34	0.16	9.3
shooting goal (soccer)	0.07	0.07	0.07	2.0
playing tennis	0.06	0.05	0.05	1.4
high kick	0.05	0.06	0.05	1.5

in different clusters. We employ hierarchical clustering (HC) methods to obtain a tree-like structure where the number of clusters can be determined post-hoc and the interpretation is richer.

Agglomerative hierarchical clustering works by merging the most similar objects in a tree-like structure from leaf nodes to the root node in a bottom-up manner. In the beginning, it creates the same number of clusters as the number of objects where each object is a cluster of its own, called as a singleton. Then, at each step, it identifies the most similar two clusters and merges them into one cluster. This process iterates until all clusters are merged into the single highest-level cluster. The order in which clusters are merged with each other, hence, results in the hierarchical tree structure.

There exist multiple methods for calculating the distance between two clusters. Given two clusters u and v , *single-linkage* uses the shortest distance between any pair of objects from u and v , *complete linkage* uses the longest distance between any pair of objects from u and v , *average-linkage* [27] uses the average distance between all pairs of objects from u and v . Given that s and t is merged to form cluster u , *weighted-linkage* [27] uses the average of distances between s and v , and t and v to find the distance between u and v . On the other hand, *Ward's* method [28] merges two clusters which minimizes the total within-cluster variance at each step. According to [29] and [30], *Ward's* and *complete-linkage* methods tends to generate a more balanced tree with approximately equal cluster sizes.

600 activity labels i, j, \dots and the pairwise distances between them D_{ij}, \dots are fed as input to the above-mentioned set of hierarchical clustering methods. In line with its goals and underlying mechanisms, the HC process is expected to reveal the hidden hierarchy of human activities. Moreover, HC algorithms utilizing different methods to calculate the distance between clusters would result in different trees which can be compared quantitatively (e.g., the extent it improves prediction accuracy at activity group levels) and qualitatively (e.g., the extent that the shared semantics within the same activity cluster is plausible).

The trees obtained by the HC method can be cut at different depths to divide the data into the desired number of clusters.




FIGURE 4. Classification accuracy of *I3D* model on ADL hierarchies.

In this case, the maximum number of clusters is 600 where each label is a cluster of its own and there is practically no clustering. The minimum number of clusters is 1 where all labels are grouped together and there is again practically no clustering. The tree can be cut to obtain k clusters where $1 \leq k \leq 600$. 2670 videos are randomly selected from *Kinetics-600* test set (excluding the 47 videos utilized earlier) and prediction of *I3D* model is compared at different levels. The accuracy at $k = 600$ gives the accuracy with respect to the individual activity labels, the accuracy at $1 < k < 600$ gives accuracy with respect to the activity groups at different levels, and $k = 1$ gives the accuracy of 100%. For instance, when $k = 10$, the activities are organized under 10 groups and accuracy is measured based on whether the label predicted by *I3D* is in the same cluster as the ground-truth label is in.

IV. RESULTS

A. OVERVIEW ON THE ACCURACY

Utilizing the two distance measures and five hierarchical clustering methods, a total of 10 different hierarchies are obtained. Figure 4 visualises the classification accuracy of *I3D* model on randomly selected 2670 videos from *Kinetics-600* test set. Each line represents an activity hierarchy obtained by a different method, x -axis shows the hierarchy levels from bottom (more detailed activities) to top (broader categories), and y -axis shows the prediction accuracy at corresponding levels of hierarchy. The predictive accuracy is higher at the upper levels of the hierarchy. This is fitting with our human intuition. It makes the machine's predictions at high levels more trustworthy. For instance, from the original *I3D* model, we only get one lengthy final output vector. Even though the original top 1 prediction from the model is wrong, the upper-level predictions have a better chance to be right.

In very detailed levels where the average cluster size is up to around 2 (i.e., when $k > 300$ approximately), most methods perform comparably. In this regime characterized by very small-sized clusters of activities, *confidence*-based distance measure slightly outperforms *lift*-based distance measure as shown by the consistently underperforming dashed lines in this region –except for *lift*-based *single-linkage* method (dashed green line) which begins to outperform all other methods at around $k = 400$. For $k < 300$, which is the more interesting and real-world case with more clustering, the two *single-linkage* methods are the only contenders with the *confidence*-based version outperforming the other once the average cluster size is above 6, that is when $k > 100$. Overall, *single-linkage* method utilizing *confidence*-based distance (hereafter called *Confidence-Single*) performs better than the others when clusters include several items (i.e., $k < 100$) or very few items (i.e., $k > 400$) and performs very close to the method with highest accuracy when $100 < k < 400$.

When Figure 4 is evaluated as a whole, we observe that a hierarchical labeling system can dramatically improve the prediction accuracy at upper levels, particularly when the activity hierarchy is discovered with an appropriate and effective clustering method. This property carries importance for prediction systems with a large number and variety of activities, particularly when the information at different levels is necessary for specific use cases. Generally, for this to be useful, the clusters should be organized in a meaningful way. It should be feasible to label or annotate clusters such that the activity labels are consistent with their respective cluster labels at upper levels.

Next, we closely examine, annotate, and discuss the activity hierarchies generated by *Confidence-Single* which has the highest overall accuracy, and by *Confidence-Ward's* which has a more balanced structure.

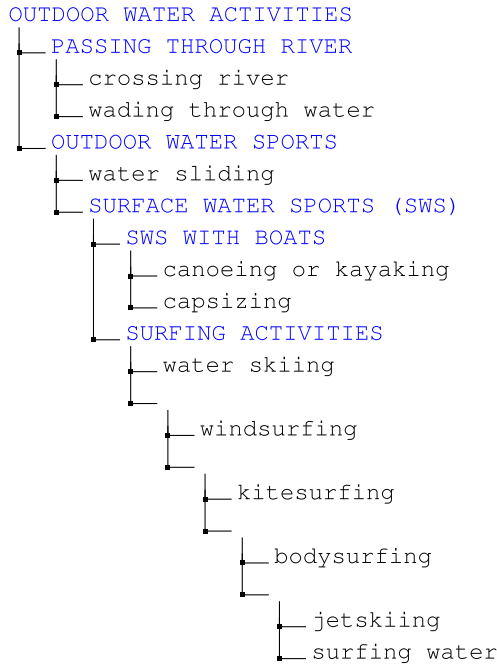


FIGURE 5. Hierarchy of outdoor water activities.




FIGURE 6. Hierarchy of hair shaping activities.

B. THE HIERARCHY OF ACTIVITIES

The activity hierarchies generated by the machine and annotated by humans cannot be produced here in full due to space limitations but will be made available online.¹ Figures 5, 6, and 7 provides some extracts from the tree where original labels are coloured in black and human annotations are written in blue and uppercase letters, and some annotations are omitted for very small clusters.

The hierarchy generated by *Confidence-Single* can be characterized as follows. The major part of the tree consists of clusters whose members are clearly related to each other. For instance, Figure 5 shows a subtree that brings together various outdoor water sports and Figure 6 shows another subtree for hair shaping activities. In general, these figures exemplify the major part where the machine is able to generate intuitive and meaningful hierarchical relations among the activity labels.

The other and minor part, however, consists of the outlier and/or under-identified activities that are not found to be sufficiently similar to most other activities. These activities,




FIGURE 7. Individual and small groups of outlier activities.

individually or in very small groups, are merged into the rest of the tree, i.e., the large tree that consists of the very intuitive groupings as described in the preceding paragraph. Figure 7 exemplifies the characteristics of this part. The singleton *riding elephant* is merged with the rest of the tree, i.e., the cluster of other 599 activities. *Playing paintball* and *playing lasertag* are grouped together to form a meaningful cluster of two before merging with the rest of the tree, i.e., a large cluster of 586 activities. Usually, activities in this part are not part of meaningful hierarchies that span more than two levels.

Given the limitations of hierarchy generated by *Confidence-Single*, we also investigate the hierarchy generated by *Confidence-Ward's* since it tends to construct clusters of comparable sizes. Unlike the former, there are no activities that are merged to the rest of the tree as individuals or in very small groups. When visually inspected, it appears as a balanced tree with some activities (those activities *Confidence-Single* is not able to place within meaningful hierarchies) are now grouped with relatively similar activities alongside some occasional misplacements. Figure 8 presents the human-annotated groups of activities based on the hierarchy automatically generated by *Confidence-Ward's*. The binary nature of the tree is omitted for neat visualization. Original activities (in *italic*) are given as examples for each cluster.

C. INSIGHTS ON THE DATASETS AND PREDICTION MODEL

The generated hierarchies provide implicit insights into the prediction model and the datasets. Some activity labels are not part of any meaningful multiple-level hierarchies and remain in very small groups or as singletons until near the end when they merge with the rest of the tree. Other activity labels are placed in hierarchies where they do not intuitively belong. This manifests under-identification for those activities. We point out three potential reasons for this problem and suggest appropriate improvements to overcome such under-identification.

First, the example videos that exhibit those activities might be scarce or monotypic in the collected datasets. This results in too few samples to learn the characteristics of these activities and their resemblance to other activities. For instance, *riding elephant* is the top prediction and in the top-5 predictions only for 4 and 19 videos, respectively. To explore whether the scarcity in the dataset is related to under-identification, consider the top 50 activities that are among the first to merge into meaningful and intuitive hierarchies

¹<https://aiiot.ucd.ie/>

- Activities involving head
 - Hair care
 - * Barber: *shaving head, getting a haircut, washing hair*
 - * Hair shaping: *combing hair, blowdrying hair, curling hair*
 - Facial gestures
 - * Eye: *winking, raising eyebrows, crossing eyes*
 - * Mouth: *yawning, laughing, sneezing*
 - Eating: *eating spaghetti, eating cake, eating watermelon*
- Food preparation: *making cheese, cooking scallops, rolling pastry, flipping pancake, chopping meat, sharpening knives*
- Playing musical instruments
 - Keyboard instruments: *playing organ, playing piano*
 - Wind instruments: *playing clarinet, playing saxophone*
 - Bowed instruments: *playing violing, playing cello*
 - Lutes: *playing ukulele, playing lute, playing guitar*
- Sports I
 - Indoor sports
 - * Gym Sports
 - Weightless: *yoga, contorting, push up*
 - Weights: *deadlifting, clean and jerk, bench pressing*
 - * Swimming: *swimming backstroke, ice swimming*
 - Outdoor Sports:
 - * Other: *tightrope walking, cartwheeling, somersaulting*
 - * Climbing: *climbing a rope, rock climbing, abseiling*
 - * In the air: *paragliding, skydiving, parasailing*
 - * Skiing: *ski jumping, skiing slalom, skiing mono*
 - * Boarding: *longboarding, skateboarding*
- Sports II
 - Outdoor water sports
 - * Involving a boat: *canoeing or kayaking, capsizing*
 - * Surfing: *kitesurfing, windsurfing, water skiing*
 - Involving a relatively large ball
 - * Football-like: *kicking soccer ball, hurling, tackling*
 - * Indoor: *playing basketball, playing volleyball*
 - Involving a relatively small ball
 - * Golf: *golf driving, golf putting, golf chipping*
 - * Bat-and-ball games: *playing cricket, hitting baseball*
 - * Throwing: *hammer throw, shot put, throwing discus*
 - Involving no ball
 - * Jumping: *triple jump, long jump, high jump*
 - * Throwing: *javelin throw*
- Other
 - Smoking and Drinking:
 - * Smoking: *smoking, smoking pipe, smooking hookah*
 - * Drinking: *tasting wine, drinking shots, bartending*
 - Handcraft: *knitting, making jewelry, bookbinding*
 - Trees: *trimming shrubs, picking fruit, trimming trees*
 - Table games: *playing monopoly, playing chess, playing poker*
 - Dancing I: *breakdancing, dancing macarena, tap dancing*
 - Dancing II: *tango dancing, swing dancing, salsa dancing*
 - Fingers: *clapping, air drumming, twiddling fingers*
 - Ice sports: *playing ice hockey, hockey stop ice skating*
 - Hitting: *punching bag, side kick, punching person (boxing)*
 - Juggling: *spinning poi, juggling balls, juggling fire*
 - ...
- * Cycling: *riding unicycle, falling off bike, jumping bicycle*

FIGURE 8. Human activity hierarchy.

and the bottom 50 activities that are merged individually or in very small groups to the rest of the tree near termination. For the top 50 activities, the median number of videos where the top five predictions include them is 117 whereas the same value for the bottom 50 videos is only 31. The overall median is 78. This indicates that the lack of videos for some activities might be preventing those activities to be placed well in the automatically generated hierarchy.

Second, the set of activities the model is trained for might not be comprehensive enough to cover other activities that are naturally very similar. For instance, consider the activity *milking cow*. The most similar activity might be *milking sheep* or some other farm activity involving *cow*. However, there is almost no other activity in the 600 activities that involve animal husbandry. As a result, *milking cow* could not be placed as part of an intuitive, interesting, or otherwise informative sub-hierarchy by *Confidence-Single*. *Confidence-Ward's* provides only a relatively meaningful hierarchy for *milking cow* where it is first merged with *grooming horse*, and then the two are merged with *bullfighting*.

Third, some activities might have very unique nature that it is hard to group them with other activities. For instance, *playing paintball* and *playing lasertag* form the very intuitive cluster *confrontational shooting sports* (in Figure 7) but they are not part of a hierarchy with meaningful depth and width since there are not many other activities neither in the model nor in the real world such that they are sufficiently close.

Tracing the individual impacts of these factors on the generated hierarchies is very difficult if not impossible. Often, more than one underlying factor contribute when some activity is not well-placed in the hierarchy. Nevertheless, our conclusion is rather simple. We recommend that in developing an image recognition model (i) data collection should be prioritized for activities that are not well-placed in the hierarchy and (ii) enhancements to the label set should be made considering the potential similarity to those activities such as including more activities from the same domain. Likewise, it is important to ensure a rich variety and quantity also for the dataset that is used in calculating the similarities between the activities for hierarchical clustering. In this way, the performance of the proposed hierarchical labeling system can be improved and become more practical for real-world tasks. Besides, the predictive accuracy might improve not only at upper levels of hierarchy but also at the level of individual activities which is desirable in any prediction use case in computer vision.

V. CONCLUSION AND FUTURE WORK

In this paper, we have proposed a method to automatically generate hierarchical activity trees to understand and improve how a HAR machine thinks. We have demonstrated that such a hierarchical labeling system can be constructed based on the similarity patterns in the prediction results from a state-of-the-art DL-based HAR model. We have performed a series of experiments to prove the feasibility of labeling systems that

can automatically generate hierarchical taxonomies of human activities from the machine's perspective.

The proposed labeling technique can preserve privacy at upper levels while maintaining a better accuracy at those desired and sufficient levels for the specific use cases. Also, deep learning video processing is generally built on one or a few datasets, with fixed label sets. Due to the nature of supervised learning, the model can only predict behaviors that exist in its training set. Even with an unseen input, a model will produce the most likely prediction from the known behaviors, which could be wrong. Through our approach, this unseen input will be at least associated with the correct higher level labels.

The proposed approach provides tools to evaluate the model's intelligence while keeping the human in the loop and reveals the bias of the underlying model. Such information can also provide guidelines for collecting training data, e.g., by focusing on the activities that are not intuitively similar but clustered together by the machine-generated tree. The contradictions in the summarized level labels between the understandings of the machine and the human, we can refer them as cognition outliers, can reveal the underlying barriers that are preventing a machine to think intelligently like a human. Future research may use not only the output from the last neural network layer but also the last few layers to further investigate the machine's presentations for the similarities and distinctions between activities. If we can understand machine thinking on a deeper level, the better we can train it to truly think like a human. Such methods would also be contributing significantly to the research stream of Explainable AI. We also motivate researchers to further develop model architectures that can directly produce predicting outputs in a hierarchical tree structure. Lastly, such machine-generated hierarchies can be used in other research problems where similar hierarchies are created manually.

REFERENCES

- [1] J. Pearl and D. Mackenzie, *The Book of Why: The New Science of Cause and Effect*, 1st ed. New York, NY, USA: Basic Books, 2018.
- [2] D. Gunning, "Explainable artificial intelligence (xai)," *Defense Adv. Res. Projects Agency (DARPA)*, vol. 2, p. 2, Nov. 2017.
- [3] A. Adadi and M. Berrada, "Peeking inside the black-box: A survey on explainable artificial intelligence (XAI)," *IEEE Access*, vol. 6, pp. 52138–52160, 2018.
- [4] A. Holzinger, C. Biemann, C. S. Pattichis, and D. B. Kell, "What do we need to build explainable AI systems for the medical domain?" 2017, *arXiv:1712.09923*. [Online]. Available: <http://arxiv.org/abs/1712.09923>
- [5] A. Yassine, A. A. N. Shirehjini, and S. Shirmohammadi, "Smart meters big data: Game theoretic model for fair data sharing in deregulated smart grids," *IEEE Access*, vol. 3, pp. 2743–2754, 2015.
- [6] I.-H. Bae, "An ontology-based approach to ADL recognition in smart homes," *Future Gener. Comput. Syst.*, vol. 33, pp. 32–41, Apr. 2014.
- [7] K. Xia, J. Huang, and H. Wang, "LSTM-CNN architecture for human activity recognition," *IEEE Access*, vol. 8, pp. 56855–56866, 2020.
- [8] L. M. Dang, K. Min, H. Wang, M. J. Piran, C. H. Lee, and H. Moon, "Sensor-based and vision-based human activity recognition: A comprehensive survey," *Pattern Recognit.*, vol. 108, Dec. 2020, Art. no. 107561.
- [9] J. Carreira and A. Zisserman, "Quo vadis, action recognition? A new model and the kinetics dataset," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jul. 2017, pp. 4724–4733.
- [10] J. M. Chaquet, E. J. Carmona, and A. Fernández-Caballero, "A survey of video datasets for human action and activity recognition," *Comput. Vis. Image Understand.*, vol. 117, no. 6, pp. 633–659, Jun. 2013.
- [11] T. Singh and D. K. Vishwakarma, "Video benchmarks of human action datasets: A review," *Artif. Intell. Rev.*, vol. 52, no. 2, pp. 1107–1154, Aug. 2019.
- [12] R. Singh, A. Sonawane, and R. Srivastava, "Recent evolution of modern datasets for human activity recognition: A deep survey," *Multimedia Syst.*, vol. 26, no. 2, pp. 83–106, Apr. 2020.
- [13] J. Carreira, E. Noland, A. Banki-Horvath, C. Hillier, and A. Zisserman, "A short note about kinetics-600," 2018, *arXiv:1808.01340*. [Online]. Available: <http://arxiv.org/abs/1808.01340>
- [14] J. Carreira, E. Noland, C. Hillier, and A. Zisserman, "A short note on the kinetics-700 human action dataset," 2019, *arXiv:1907.06987*. [Online]. Available: <http://arxiv.org/abs/1907.06987>
- [15] H. Kuehne, H. Jhuang, E. Garrote, T. Poggio, and T. Serre, "HMDB: A large video database for human motion recognition," in *Proc. Int. Conf. Comput. Vis. (ICCV)*, Nov. 2011, pp. 2556–2563.
- [16] K. Soomro, A. R. Zamir, and M. Shah, "UCF101: A dataset of 101 human actions classes from videos in the wild," 2012, *arXiv:1212.0402*. [Online]. Available: <http://arxiv.org/abs/1212.0402>
- [17] F. C. Heilbron, V. Escorcia, B. Ghanem, and J. C. Niebles, "ActivityNet: A large-scale video benchmark for human activity understanding," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2015, pp. 961–970.
- [18] Y.-G. Jiang, Z. Wu, J. Wang, X. Xue, and S.-F. Chang, "Exploiting feature and class relationships in video categorization with regularized deep neural networks," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 40, no. 2, pp. 352–364, Feb. 2018.
- [19] Y. Tang, D. Ding, Y. Rao, Y. Zheng, D. Zhang, L. Zhao, J. Lu, and J. Zhou, "COIN: A large-scale dataset for comprehensive instructional video analysis," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2019, pp. 1207–1216.
- [20] P. Wozniowski, E. Tonkin, and P. Flach, "Activities of daily living ontology for ubiquitous systems: Development and evaluation," *Sensors*, vol. 18, no. 7, p. 2361, Jul. 2018.
- [21] T. Long, P. Mettes, H. T. Shen, and C. G. M. Snoek, "Searching for actions on the hyperbole," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2020, pp. 1138–1147.
- [22] A. Gaidon, Z. Harchaoui, and C. Schmid, "Activity representation with motion hierarchies," *Int. J. Comput. Vis.*, vol. 107, no. 3, pp. 219–238, May 2014.
- [23] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "ImageNet: A large-scale hierarchical image database," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2009, pp. 248–255.
- [24] R. Agrawal, T. Imieliński, and A. Swami, "Mining association rules between sets of items in large databases," in *Proc. ACM SIGMOD Int. Conf. Manage. Data (SIGMOD)*, 1993, pp. 207–216.
- [25] R. J. Bayard and R. Agrawal, "Mining the most interesting rules," in *Proc. 5th ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining (KDD)*, 1999, p. 145–154.
- [26] J. Han, J. Pei, and M. Kamber, *Data Mining: Concepts and Techniques*. Amsterdam, The Netherlands: Elsevier, 2011.
- [27] R. Sokal and C. Michener, "A statistical method for evaluating systematic relationships," *Univ. Kansas Sci. Bull.*, vol. 38, pp. 1409–1438, Mar. 1958.
- [28] J. H. Ward, "Hierarchical grouping to optimize an objective function," *J. Amer. Stat. Assoc.*, vol. 58, no. 301, pp. 236–244, Mar. 1963.
- [29] S. Hands and B. Everitt, "A Monte Carlo study of the recovery of cluster structure in binary data by hierarchical clustering techniques," *Multivariate Behav. Res.*, vol. 22, no. 2, pp. 235–243, Apr. 1987.
- [30] N. H. Timm, *Applied Multivariate Analysis: Springer Texts in Statistics*. New York, NY, USA: Springer-Verlag, 2002.



MAHSUN ALTIN received the B.Sc. degree in computer engineering from Abdullah Gül University, Turkey, and completed an internship program with University College Dublin, Ireland. He has completed a project with an International Team, NOKIA, for an Erasmus project. He is currently working as a Data Scientist with TurkNet Telecommunication Company, Turkey, focusing on big data analysis by using advanced analysis tools and machine learning algorithms.



FURKAN GÜRSOY received the B.Sc. degree in management information systems from Boğaziçi University, Turkey, and the M.Sc. degree in data science from İstanbul Şehir University, Turkey. He is currently pursuing the Ph.D. degree in management information systems with Boğaziçi University. He previously worked as an IS/IT Consultant and a Machine Learning Engineer with the industry for several years. He held a Visiting Researcher Position with IMT Atlantique, France,

in 2020. His research interests include complex networks, machine learning, simulation, and broad data science.



LINA XU received the B.E. degree from the Software Engineering School, Fudan University, China, the B.Sc. degree in computer science from University College Dublin (UCD), Ireland, through a Fudan and UCD joint Program, and the Ph.D. degree from UCD, in 2014. She worked with HP Labs as a Research Scientist from the end of 2014 to 2016. She is currently working as an Assistant Professor with UCD. She is also working on a smart living project aiming to apply

IoT technologies to smart living. Her research interests include the Internet of Things (IoT), 5G networks, smart networking, and machine learning.

• • •