From Actions to Events: A Transfer Learning Approach Using Improved Deep Belief Networks

Mateus Roder*, Jurandy Almeida†, Gustavo H. de Rosa*, Leandro A. Passos*,
André L. D. Rossi*, João P. Papa*

*Department of Computing, São Paulo State University – UNESP, Bauru, Brazil
{mateus.roder, gustavo.rosa, leandro.passos, andre.rossi, joao.papa}@unesp.br
†Instituto de Ciência e Tecnologia, Universidade Federal de São Paulo – UNIFESP, São José dos Campos, Brazil
jurandy.almeida@unifesp.br

Abstract—In the last decade, exponential data growth supplied machine learning-based algorithms' capacity and enabled their usage in daily-life activities. Additionally, such an improvement is partially explained due to the advent of deep learning techniques, i.e., stacks of simple architectures that end up in more complex models. Although both factors produce outstanding results, they also pose drawbacks regarding the learning process as training complex models over large datasets are expensive and time-consuming. Such a problem is even more evident when dealing with video analysis. Some works have considered transfer learning or domain adaptation, i.e., approaches that map the knowledge from one domain to another, to ease the training burden, yet most of them operate over individual or small blocks of frames. This paper proposes a novel approach to map the knowledge from action recognition to event recognition using an energy-based model, denoted as Spectral Deep Belief Network. Such a model can process all frames simultaneously, carrying spatial and temporal information through the learning process. The experimental results conducted over two public video dataset, the HMDB-51 and the UCF-101, depict the effectiveness of the proposed model and its reduced computational burden when compared to traditional energy-based models, such as Restricted Boltzmann Machines and Deep Belief Networks.

I. INTRODUCTION

Machine Learning (ML) techniques emerged in the last decades as revolutionary tools capable of solving or slightly alleviating the burden imposed by repetitive and tedious tasks. Recently, the advent of Deep Learning (DL) algorithms powered up such advances, providing astonishing predictions over complex domains. On the other hand, video-based domain tasks still pose challenging assignments to intelligent algorithms, e.g., recognizing human actions in videos [1].

An action recognition task is characterized by an observation of a complete sequence of movements performed by a human followed by its classification [2]. Such a task plays a fundamental role in video surveillance-based security systems, content-based video retrieval, and self-driving cars, among others. Event recognition is a similar task which models the ability to retrieve specific actions from sequences of videos, focusing on learning behaviors related to events of interest, i.e., events which comprise specific activities and objects in a given scene. Thus, there is a slight difference between action and event recognition tasks, i.e., while the former attempts to identify any action performed in the scene, the latter remarks specific movements, such as "is there anybody drinking in the street?" or "does this video contains any unusual behavior?". Event recognition approaches are commonly employed for monitoring public or private areas in search for anomalous behaviors, such as violent assaults on sports events, abandoned objects on train stations, and route obstruction on industrial environments, among others.

Despite their similarities, both action and event recognition tasks present their particularities, posing distinct challenges while training an intelligent model. An interesting approach used to solve event recognition tasks is importing some correlated knowledge extracted over action recognition-based models, denoted as transfer learning.

Transfer learning studies the possibilities of transferring knowledge from source domains to different contexts (target domains). To illustrate such an idea, consider an autonomous-driving car trained with road-traffic data from New Zealand. It will not work effectively on Brazilian streets due to different signage and road-traffic rules, a distinct influx of vehicles, and right-lane-based driving, among other issues. However, adapting the knowledge learned in New Zealand to Brazil can reduce computational costs and time needed to train new models [3]. In a nutshell, the transfer learning considers two or more feature space distributions been related or equal, where the task related to each one can be also related or equal, providing useful information for the target task [4].

Several works addressed the problem of action and event recognition through transfer learning strategies. Farajidavar et al. [5], for instance, proposed a transductive transfer learning method for action recognition in tennis games. Further, Shao et al. [6] published a survey comprising several approaches for these tasks, such as space approximation [7], Gaussian mixture model [8], and geometric reasoning [9]. Recently, novel approaches considered DL methods, such as Wang et al. [10], who introduced Generative Adversarial Networks for action recognition using partial-modalities. Tas et al. [11] employed a Convolutional Neural Network (CNN) for action recognition and supervised domain adaptation on 3D body skeletons. Furthermore, Gao et al. [12] introduced a two-stream graph CNN for zero-shot action recognition, while Liu et al. [13] proposed image-to-video adaptation and fusion networks in the same context.

Among several DL algorithms, an energy-based model

known as Deep Belief Network (DBN) [14] obtained considerable popularity in the last years due to its notable results in a wide variety of applications [15], [16]. It is composed of multiple hidden layers, such that each layer is a greedily trained Restricted Boltzmann Machine (RBMs) [17]. Even though some works proposed DBNs for action [18], [19] and a specific type of event recognition [20], as far as we know, there is still no work that has successfully addressed the concepts of transfer learning from actions to events with video through DBNs, i.e., to learn useful features from highly structured actions and movements to the generalization of high-level events. Therefore, the main contributions of this work are threefold: (i) to introduce DBNs in the video event classification domain, (ii) to propose two approaches, denoted as Aggregative-DBN and Gradient-DBN, that employ frame fusion and image gradient respectively, and (iii) to support the lack of video-based event recognition in literature.

Additionally, we consider the following hypotheses: (i) DBNs are able to learn useful correlations that map actions to high-level events in video-based domain tasks; (ii) the overall accuracy can be improved using the proposed approaches; and (iii) the overall training time can be reduced with the Aggregative-based approach.

The remainder of this paper is organized as follows. Section II introduces the main theoretical concepts used in the manuscript, while Section III presents the proposed approaches. Further, Sections IV and V describe the methodology adopted in this work and the experimental results. Finally, Section VI states conclusions and future work.

II. THEORETICAL BACKGROUND

A. Video-based Domain

Let $\mathcal{F} = \{F_1, F_2, \ldots, F_n\}$ be a temporal sequence of frames (images) F_i, which represents a soundless scene and possible movements of its components. Such frames can be classified according to the complexity of their internal representations and the interaction level between their entities. Thus, it is possible to generate four classification-based categories: attributes and movements, low-level events and actions, interaction, and high-level events [21].

Figure 1 illustrates the hierarchical form of the categories discussed hereafter. The movement characterizes the lowest representation level of a frame and is widely employed to recognize human actions, such as body movement [22]. On the other hand, low-level events and actions represent a particular chain of movements, usually carried out by an entity, e.g., car or person. Additionally, if such actions are performed or interacted by more than one entity or object, it is possible to categorize them as an interaction [21]. Finally, the highest-level category, denoted as complex events, represents the interaction of entities or sequence of actions in a specific time window in the video. For instance, one can identify a birthday party as an event composed of several actions and entities in a single scene.

Therefore, the event recognition task attempts to detect the complex events' spatial and temporal locations in a sequence

of frames [21]. The literature lacks in standardizing the difference between actions and events, making them interchangeable in most applications [21].
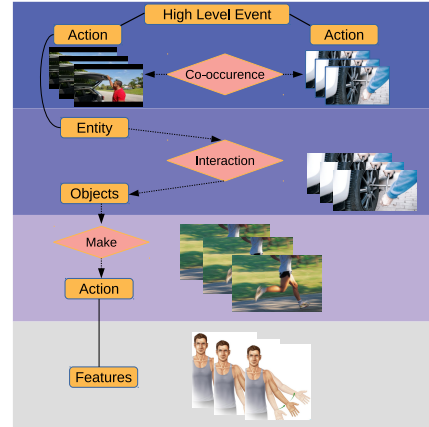

Fig. 1. Video-based domain hierarchical complexity.

B. Applied Transfer Learning

Transfer learning has recently received attention due to the advent of the ImageNet[1] dataset and the increased processing power of GPUs. Such a task consists in transferring the knowledge from the source domain to the target domain, providing information that can be useful on the target domain, mainly when data are insufficient or the computational resources are scarce. In this way, it is possible to train deep neural networks in large-scale image/video datasets and use them to fine-tune more specific tasks [23].

Given the previous concepts, it is possible to elucidate the mathematical formulation regarding the problem addressed here. Let Γ be a high-level event recognition task, as well as let $\mathcal{D}_S$ and $\mathcal{D}_T$ be the source (action domain) and the target space domains (event domain), respectively. Additionally, the source domain is composed of the subspaces $\mathcal{A} \in \mathbb{R}^{d_a}$, $\mathcal{M} \in \mathbb{R}^{d_m}$, and $\mathcal{I} \in \mathbb{R}^{d_i}$, where $\{\mathcal{A}, \mathcal{M}, \mathcal{I}\} \subset \mathcal{D}_S$, while the target domain is composed of $\mathcal{E} \in \mathbb{R}^{d_e}$, where $\{\mathcal{E}\} \subset \mathcal{D}_T$. The subspace $\mathcal{A}$ stands for the d_a-dimensional base actions, $\mathcal{M}$ stands for the d_m-dimensional movements, $\mathcal{I}$ represents the interactions between d_i-dimensional entities, and $\mathcal{E}$ stands for the d_e-dimensional high-level events.

From the transfer learning theory [24], a specific case is when the data from the source domain and target domain keep their probabilities, letting $\mathcal{D}_S$ to be identical to $\mathcal{D}_T$, differing only on the target task, which is the problem addressed in this paper. Finally, it is possible to formulate the proposed approach for the given task using Equation 1, as follows:

$$\Gamma = \{y_{\mathcal{D}_T}, f(\mathcal{D}_S)\}, \tag{1}$$

where $y_{\mathcal{D}_T}$ stands for the target domain labels and $f(\mathcal{D}_S)$ for the function that learns features from the source domain. Such learned features are useful for the target domain and

---

[1]https://image-net.org/

its respective event classification task, i.e., a neural network that learns from $\mathcal{D}_S$ and is fine-tuned in $\mathcal{D}_T$ with labels $y_{\mathcal{D}_T}$. Here, $\mathcal{D}_T$ has the same probability distribution that $\mathcal{D}_S$, as aforementioned.

C. Restricted Boltzmann Machines

Restricted Boltzmann Machines (RBMs) [17], [25] are described as a bipartite graph composed of two layers of neurons, i.e., a visible layer $\mathbf{v} \in \{0, 1\}^m$, which is responsible for the input data, and a hidden layer $\mathbf{h} \in \{0, 1\}^n$, whose units map the data representation into a latent space. This interaction is modeled by a weight matrix, $\mathbf{W} \in \Re^{m \times n}$, which connects each visible unit v_i to all hidden units h_j, and vice-versa, denoted by the arc w_{ij}.

The RBM learning procedure is performed by the minimization of an energy function concerning some intrinsic variables, described as follows:

$$E(\mathbf{v}, \mathbf{h}) = -\sum_{i=1}^{m} b_i v_i - \sum_{j=1}^{n} c_j h_j - \sum_{i=1}^{m} \sum_{j=1}^{n} v_i h_j w_{ij}, \quad (2)$$

where $\mathbf{b} \in \Re^m$ and $\mathbf{c} \in \Re^n$ stand for the bias vector considering the visible and hidden layers, respectively.

Computing the joint probability of the system poses an intractable task due to the increasing number of possible states. However, since the model is represented as a bipartite graph, one can compute both the visible and hidden units' activation in a mutually independent fashion, performed as follows:

$$P(v_i = 1|\mathbf{h}) = \phi \left(\sum_{j=1}^{n} w_{ij} h_j + b_i \right), \quad (3)$$

and

$$P(h_j = 1|\mathbf{v}) = \phi \left(\sum_{i=1}^{m} w_{ij} v_i + c_j \right). \quad (4)$$

Note that $\phi(\cdot)$ stands for the logistic-sigmoid function. Finally, we can solve the equations above by iteratively sampling over a Markov Chain, using the well-known Contrastive Divergence (CD) algorithm [17].

The learning process in RBMs consists of an optimization problem whose goal is to minimize the energy function given in Equation 2. In other words, such a process ends up maximizing the marginal probability distribution of the visible units, defined as follows:

$$P(\mathbf{v}) = \frac{\sum_{\mathbf{h}} e^{-E(\mathbf{v}, \mathbf{h})}}{\sum_{\mathbf{v}, \mathbf{h}} e^{-E(\mathbf{v}, \mathbf{h})}}, \quad (5)$$

which is commonly handled in its natural logarithm version, i.e., more precisely, we aim at maximizing the negative logarithm of the likelihood function (Negative Log-Likelihood - NLL). Moreover, regarding the visible units, such a procedure

can be easily extended to the continuous domain, which is useful to model any type of input. The changes occur on the energy function, as follows:

$$E(\mathbf{v}, \mathbf{h}) = \sum_{i=1}^{m} \frac{(v_i - b_i)^2}{2\sigma_i^2} - \sum_{j=1}^{n} c_j h_j - \sum_{i=1}^{m} \sum_{j=1}^{n} \frac{v_i}{\sigma_i} h_j w_{ij}. \quad (6)$$

Considering the derivatives, it is straightforward to show that the visible prior becomes:

$$P(v_i = 1|\mathbf{h}) \sim \mathcal{N} \left(\sum_{j=1}^{n} w_{ij} h_j + b_i, \sigma_i^2 \right), \quad (7)$$

and, when a Gaussian input with zero mean and one unit standard deviation is employed, such a prior becomes simple and easy to sample, since the whole procedure is still the same.

D. Deep Belief Networks

Deep Belief Networks (DBNs) [14] are generative graphical models composed of a stack of RBMs, thus providing multiple layers of latent variables, such that the hidden layer of the bottommost RBM is employed to feed the subsequent input units successively until reaching the topmost layer. DBNs are trained greedily, meaning that an RBM at a specific layer does not consider others during its learning procedure. Thus, a DBN is composed of L layers, where $\mathbf{W}^l$ is the weight matrix of an RBM at layer l. Additionally, we can observe that the hidden units at layer l become the input units of layer $l + 1$.

The aforementioned procedure stands for the generative pre-training. Afterward, it is possible to attach fully-connected (FC) layers with softmax outputs at the topmost hidden layer for a discriminative fine-tuning, which can be used for classification tasks.

III. PROPOSED APPROACHES

This work proposes to employ DBNs as non-linear functions $f(\mathcal{D}_S)$ to learn from the source domain, $\mathcal{D}_S$, the information that can be used to map the target domain, $\mathcal{D}_T$, i.e., to extract information from videos and use them to classify high-level events. Additionally, two approaches are shown to the respective task.

A. Aggregative Deep Belief Networks

The first alternative architecture is denoted as Aggregative (A- prefix in models), which modifies the DBNs' first layer. Such a variation is designed considering two main concepts: (i) to capture general spatio-temporal information without additional techniques, such as optical flow algorithms, and (ii) reduce the overall computational burden.

The proposed A-DBN processes all frames simultaneously instead of processing one frame at a time, enabling a complete parameter update at each iteration. In other words, A-DBN aggregates all frames $\{F_1, F_2, \ldots, F_n\}$ into a single frame denoted as F_r, which represents their summation. This procedure carries spatial information (contour and edges) along their temporal trajectories, highlighted as "spectrum" in the

resulting frame. Figure 2 depicts the F_r aggregation process, as well as its highlighted region. Afterward, A-DBN first layer infers the posterior distribution given F_r, as follows:

$$P(h_j = 1|\boldsymbol{F_r}) = \phi\left(\sum_{i=1}^{m} w_{ij}F_{ri} + c_j\right). \qquad (8)$$


Fig. 2. Frames aggregation for the Aggregative-based approach.

B. Gradient Deep Belief Networks

The second alternative architecture denoted as Gradient (G-prefix in models) also modifies the DBNs first layer. Such a variation is designed considering one main concept, i.e., to capture general motion information between two frames.

The proposed G-DBN processes two consecutive frames instead of processing one frame at a time. In other words, G-DBN defines a resultant frame, F_r, as the direct subtraction as follows: given F_1 and F_2, F_r stands for $F_2 - F_1$. This procedure carries motion cues from the spatial domain along their trajectories. Figure 3 depicts the F_r generation process, as well as its highlighted region. Afterward, G-DBN first layer infers the posterior distribution given F_r, as follows:

$$P(h_j = 1|\boldsymbol{F_r}) = \phi\left(\sum_{i=1}^{m} w_{ij}F_{ri} + c_j\right). \qquad (9)$$


Fig. 3. Frames generation for the Gradient-based approach.

IV. Experiments

In this section, we describe the dataset and the experimental setup employed to apply and compare DBNs with A-DBN, proposed in this paper, for the task of event recognition from complex actions domain.

A. Dataset

We opted to use two well-known datasets, UCF-101 [26] and HMDB-51 [27], as they represent a big challenge and are well-established video action recognition datasets. Both datasets comprises a significant amount of data from real-world action videos collected from YouTube and classified among 101 and 51 distinct classes, respectively. Moreover, such a diversity becomes more expressive due to the substantial variations in camera motion, object appearance and pose, scaling, viewpoint, cluttered background, and illumination conditions.

The 13,320 videos from UCF-101 are grouped in 5 macro-categories, which are easily interpreted as high-level events. Also, this mapping expects inter-class videos to share standard and essential features, which helps in recognizing common actions and interactions. Following the authors' guideline, the high-level events used: (0) sports practice; (1) musical practice with an instrument; (2) human-object interaction; (3) human body-motion; and (4) people interacting. Random clips depicting such classes are presented in Figure 4, each class is represented by the color of the border: green for 0, light-blue for 1, blue for 2, red for 3, and purple for 4.

Fig. 4. Random clips from UCF-101 [26].

Similarly, the 6,766 videos from HMDB-51 are grouped in 5 macro-categories, easily interpreted as high-level events. Also, following the authors' guideline, the high-level events are (0) human facial expression; (1) manipulation of objects in the face region; (2) body movement; (3) interaction between people and object(s); and (4) person interacting with each other, where numbers in parentheses represent classes. Additionally, frames of clips from the dataset are shown in Figure 5, where the color of the border represents the event class: green for 0, light blue for 1, blue for 2, red for 3, and purple for 4.

Both datasets provide three partitions, each with separate data for training and testing, where the first partition was used in this study, as it seems to have the most difficult test samples to evaluate as cluttered background or fewer interactions and actions. The process of splitting/acquiring the frames is performed in a similar way to the work of Ng et al. [28], using 6 frames per video clip uniformly distributed over time. In their work, Ng et al. showed that 6 frames per video are enough to ensure a good performance, achieving the same results as 20 frames, for example, in addition to imprinting a lower computational load on the action classification task.

Regarding the pre-processing step, two transformations

Fig. 5. Random clips from HMDB-51 [27].

were employed before the image conversion to grayscale. The first transformation concerns cutting operations, removing black regions that do not carry information, and resizing from the original size (240×320) to 72×96, to facilitate the processing of energy-based models. The second transform stands for feature-normalization using a Gaussian distribution with zero mean and unit variance.

B. Experimental Setup

Regarding the experimental setup hardware, we employed an Intel 2x Xeon(R) E5-2620 @ 2.20GHz (40 cores), a GTX 1080 Ti, and 128 GB of RAM. For the unsupervised pre-training process, we opted to use mini-batches of 128 samples and 3 epochs per-layer. Finally, Table I describes the employed architectures and hyper-parameters.

Each model is connected to two additional fully-connected layers that are fine-tuned using the well-known Adam optimizer [29] with the learning rate equals to 10^{-3} and the same number of epochs and mini-batch size. Such FC layers have two configurations since they depend on the number of hidden neurons from the last RBM/DBN layer, i.e., $2,000 - 1,000 - 5$ and $4,000 - 2,000 - 5$. It is important to highlight that the RBMs were employed for $2,000$ hidden neurons only as this work primarily focuses on using the hierarchical information learned by DBNs.

Furthermore, we opted to use the same approach employed by transfer learning, i.e., to freeze the connections from the first hidden layer and make a gentle adjustment of subsequent hidden layers (learning rate equals to 10^{-6}). The cross-entropy loss was used when adjusting model weights during the fine-tuning process, while the final measure was the accuracy on the testing set. Finally, to mitigate any stochastic nature, each model was fully trained and fine-tuned for 6 repetitions.

V. Experimental Results

This section presents the experimental results concerning the DBNs and the proposed approaches, i.e., the A-DBN and the G-DBN, applied to the task of event recognition. All models follow the work of Ng et al. [28], using 6 frames per video clip uniformly distributed over time.

A. Model Evaluation for Event Recognition

Regarding the main task, i.e., learning from actions to classify high-level events, Tables II and III show the predictive performance for all models and architectures. The highlighted result stands for the best mean accuracy. Besides, it also presents the average running time for each model, averaged over the six repetitions to analyze the models' computation impact and efficiency.

From Table II, we can notice prominent results, mainly for the proposed Aggregative version of DBNs. Starting from the base model, A-RBM achieved a better accuracy than RBM, i.e., 42.48% against 38.71%, representing a meaningful difference (3.77 points of mean percentual accuracy), while the G-RBM model achieved 44.04%, outperforming also the Aggregative approach. In addition, A-RBM has a significantly lower computational burden, approximately 14% less running time. To clarify, hereafter, the percentual difference between models' accuracy stands for the absolute mean value of the proposed model (A- or G-) subtracted to its standard version, and for the running time, it is the mean proposed approach time divided by its standard model.

Regarding the second architecture, i.e., the α models, the Aggregative version overcomes the standard DBN and G-DBN in mean accuracy by almost 7% and 4.5%, respectively, representing a meaningful improvement. However, the standard model and the G-DBN do not overpass its parameterless version (RBM and G-RBM) in mean accuracy. Moreover, the running time for A-DBN$_\alpha$ was 30% smaller than DBN$_\alpha$, impacting positively to a lighter training burden.

Concerning the β models, a similar behavior was observed, i.e., the Aggregative version overpassed the standard DBN in approximately 4% in mean accuracy, and almost 33% less running time. Also, the G-DBN model achieved the best mean accuracy, 44.84%, with a small standard deviation. However, such results show that adding more hidden layers does not lead to an impressive performance improvement for A-DBN since its mean accuracy was close to the A-DBN$_\alpha$, while for the G-DBN$_\beta$ the performance was increased. Moreover, even with the previous observation, the A-DBN$_\beta$ still gives a better running time than its baselines.

Regarding the fourth architecture (ι models), the same behavior was observed, highlighting that A-DBN$_\iota$ achieved a remarkable mean accuracy of 45.01%, the highest average value over the baselines. Moreover, the running time for A-DBN$_\iota$ was 30% better than its standard version. Here, the Aggregative models showed that more hidden units might benefit the overall performance for the event classification task with transfer learning.

Finally, the ζ models showed almost the same results as the ι models, mainly for the proposed approaches, A-DBN and G-DBN, which achieved a mean accuracy of 44.87% and 44.86%, respectively. Such results indicate that more hidden neurons improve the models' performance. However, the larger versions may demand more epochs of pre-training and/or data. Also, the running time of A-DBN overpasses its standard and

Model	Layers	Hidden Neurons	Momentum	Learning Rate
RBM	1	$[2,000]$	$[0.5]$	$[1 \cdot 10^{-3}]$
A-RBM	1	$[2,000]$	$[0.5]$	$[1 \cdot 10^{-3}]$
G-RBM	1	$[2,000]$	$[0.5]$	$[1 \cdot 10^{-3}]$
DBN$_\alpha$	2	$[2,000 - 2,000]$	$[0.5; 0.5]$	$[1 \cdot 10^{-3}; 5 \cdot 10^{-4}]$
DBN$_\beta$	3	$[2,000 - 2,000 - 2,000]$	$[0.5; 0.5; 0.5]$	$[1 \cdot 10^{-3}; 5 \cdot 10^{-4}; 5 \cdot 10^{-4}]$
DBN$_\iota$	2	$[4,000 - 4,000]$	$[0.5; 0.5]$	$[5 \cdot 10^{-4}; 5 \cdot 10^{-4}]$
DBN$_\zeta$	3	$[4,000 - 4,000 - 4,000]$	$[0.5; 0.5; 0.5]$	$[5 \cdot 10^{-4}; 5 \cdot 10^{-4}; 5 \cdot 10^{-4}]$
A-DBN$_\alpha$	2	$[2,000 - 2,000]$	$[0.5; 0.5]$	$[1 \cdot 10^{-3}; 5 \cdot 10^{-4}]$
A-DBN$_\beta$	3	$[2,000 - 2,000 - 2,000]$	$[0.5; 0.5; 0.5]$	$[1 \cdot 10^{-3}; 5 \cdot 10^{-4}; 5 \cdot 10^{-4}]$
A-DBN$_\iota$	2	$[4,000 - 4,000]$	$[0.5; 0.5]$	$[5 \cdot 10^{-4}; 5 \cdot 10^{-4}]$
A-DBN$_\zeta$	3	$[4,000 - 4,000 - 4,000]$	$[0.5; 0.5; 0.5]$	$[5 \cdot 10^{-4}; 5 \cdot 10^{-4}; 5 \cdot 10^{-4}]$
G-DBN$_\alpha$	2	$[2,000 - 2,000]$	$[0.5; 0.5]$	$[1 \cdot 10^{-3}; 5 \cdot 10^{-4}]$
G-DBN$_\beta$	3	$[2,000 - 2,000 - 2,000]$	$[0.5; 0.5; 0.5]$	$[1 \cdot 10^{-3}; 5 \cdot 10^{-4}; 5 \cdot 10^{-4}]$
G-DBN$_\iota$	2	$[4,000 - 4,000]$	$[0.5; 0.5]$	$[5 \cdot 10^{-4}; 5 \cdot 10^{-4}]$
G-DBN$_\zeta$	3	$[4,000 - 4,000 - 4,000]$	$[0.5; 0.5; 0.5]$	$[5 \cdot 10^{-4}; 5 \cdot 10^{-4}; 5 \cdot 10^{-4}]$

Architecture	Accuracy	Time
RBM	38.71 ± 1.04	315.00 ± 5.00
A-RBM	42.48 ± 0.94	270.00 ± 5.00
G-RBM	44.04 ± 1.82	314.00 ± 5.00
DBN$_\alpha$	37.72 ± 4.67	765.00 ± 5.00
A-DBN$_\alpha$	44.66 ± 1.28	540.00 ± 5.00
G-DBN$_\alpha$	40.16 ± 11.63	764.00 ± 5.00
DBN$_\beta$	40.55 ± 3.54	$1,215.00 \pm 5.00$
A-DBN$_\beta$	44.80 ± 2.02	810.00 ± 5.00
G-DBN$_\beta$	44.84 ± 0.08	$1,211.00 \pm 5.00$
DBN$_\iota$	41.92 ± 2.65	775.00 ± 6.00
A-DBN$_\iota$	$\mathbf{45.01 \pm 1.39}$	550.00 ± 6.00
G-DBN$_\iota$	44.84 ± 0.04	773.00 ± 6.00
DBN$_\zeta$	42.33 ± 4.51	$1,225.00 \pm 6.00$
A-DBN$_\zeta$	44.87 ± 2.81	820.00 ± 6.00
G-DBN$_\zeta$	44.86 ± 0.14	$1,220.00 \pm 6.00$

on mean accuracy due to one specific run that pushed down the performance (note the standard deviation). Overall, the performance improvement observed can be directly linked to the higher abstraction achieved by hidden layers. Such approaches can improve the DBN lower bound and provide a further improvement in discriminative fine-tuning. Besides, they were also pre-trained with a relatively small number of epochs, which can induce a less efficient overall lower bound optimization. However, the results showed that more hidden units in hidden layers improve the mean accuracy rate, as the A-DBN$_\iota$ and A-DBN$_\zeta$ models have shown.

From Table III, one can notice interesting results, mainly for the proposed approaches. The base model, A-RBM achieved a better accuracy rate than RBM, i.e., 35.19% against 34.60%, representing a meaningful difference (0.60 points of mean percentual accuracy), while the G-RBM model achieved 38.05%, beating also the Aggregative approach. Here, it is interesting to note that the time difference between models was not so impressive, explained by the low data volume.

Regarding the second architecture, i.e., the α models, the Aggregative version overcomes the standard DBN in mean accuracy by almost 3%, while the G-DBN overpassed its standard version in approximately 4%, respectively, representing a meaningful improvement. Moreover, it is important to notice that the running time for A-DBN$_\alpha$ was 8% smaller than DBN$_\alpha$, impacting positively to a lighter training burden.

Concerning the β models, was observed a similar behavior, i.e., the Aggregative version overpassed the standard DBN in approximately 5% in mean accuracy, and almost 10% less running time. Also, the G-DBN model overpassed its standard version in approximately 5% of mean accuracy, with a lower

the Gradient version by approximately 33%.

Nonetheless, it is essential to notice that A-DBN models have no difficulty in overpassing the A-RBM mean accuracy; however, the G-DBN$_\alpha$ do not overpass its simpler version

Architecture	Accuracy	Time
RBM	34.60 ± 3.90	45.00 ± 5.00
A-RBM	35.19 ± 4.24	30.00 ± 5.00
G-RBM	38.05 ± 0.36	44.00 ± 5.00
DBN_α	34.49 ± 3.98	144.00 ± 5.00
$A\text{-}DBN_\alpha$	37.68 ± 4.19	132.00 ± 5.00
$G\text{-}DBN_\alpha$	38.40 ± 0.01	143.00 ± 5.00
DBN_β	33.83 ± 4.06	216.00 ± 5.00
$A\text{-}DBN_\beta$	38.70 ± 4.33	195.00 ± 5.00
$G\text{-}DBN_\beta$	38.40 ± 0.01	214.00 ± 5.00
DBN_ι	34.41 ± 4.03	150.00 ± 6.00
$A\text{-}DBN_\iota$	38.23 ± 4.25	138.00 ± 6.00
$G\text{-}DBN_\iota$	38.40 ± 0.01	148.00 ± 6.00
DBN_ζ	34.53 ± 4.04	225.00 ± 6.00
$A\text{-}DBN_\zeta$	$\mathbf{38.86 \pm 4.26}$	207.00 ± 6.00
$G\text{-}DBN_\zeta$	37.41 ± 2.43	223.00 ± 6.00

standard deviation. However, such results show that adding more hidden layers does not cause an impressive performance improvement for G-DBN since its mean accuracy was close to the G-DBN$_\alpha$, while for the A-DBN$_\beta$ the performance was increased. Moreover, even with the previous observation, the A-DBN$_\beta$ still gives a better running time than its baselines.

Regarding the ι models, the baseline model slightly increased its performance, however, the G-DBN achieved the highest mean accuracy, 38.40%, overpassing the DBN and A-DBN models. The Aggregative version was also better than DBN, with 38.23%, and 8% less running time than DBN. However, such results show that adding more hidden layers does not cause an impressive performance improvement for the proposed approach, keeping the hidden neurons in $2,000$ units, since the mean accuracies were close to the A-DBN$_\beta$ and G-DBN$_\beta$.

Finally, the ζ models showed an interesting performance improvement regarding the Aggregative approach, which achieved a remarkable mean accuracy of 38.86%. On the other hand, G-DBN suffers from a performance decreasing, achieving 37.41% of mean accuracy. This fact pointed out that more hidden neurons improve the Aggregative's performance while keeping the lowest mean time for training (207 minutes).

In general, one can observe two interesting behaviors for the proposed approaches, the Aggregative-based models were able to improve the models' performance, resulted from the total frames aggregation that carries general motion information. On the other hand, the Gradient-based models did not improve their performance as the Aggregative-based, explained by the fact that two-by-two frames on the employed datasets may not carry as much information like the overall sum of the six frames.

VI. CONCLUSIONS AND FUTURE WORKS

In this paper, we addressed a transfer learning approach on two well-known video datasets, learning from the action- to the event-based domain, through energy-based models, such as RBMs and DBNs. Furthermore, we proposed Aggregative-based and Gradient-based approaches that modify the processing of the frames, simplifying the model's complexity and giving robustness, saving processing time, and improving the model's generalization. Experimental results show that the proposed approach can reduce the computational time by as much as 33% regarding the A-models.

The results were promising since most A-DBN architectures achieved a feasible mean accuracy rate, highlighting the models A-DBN$_\iota$ and A-DBN$_\zeta$. Also, the Aggregative models showed a meaningful reduction in running time during the unsupervised pre-training phase. The Gradient models showed a stable behavior, varying almost nothing regarding the different architectures. Therefore, these experimental results support our hypothesis that it is possible to transfer the knowledge from actions to events with the employed energy-based models, without complex inputs such as optical flow or convolutions. Finally, one can highlight that by increasing the number of hidden neurons the overall performance improved, pointing out the models' opportunity to extract more information.

Regarding future works, we plan to investigate the effect of combining convolution operators in energy-based models, such as the Convolutional Restricted Boltzmann Machine (CRBM). Additionally, we aim at employing more complex models, such as Deep Boltzmann Machines (DBMs), trained with a more expressive number of epochs. Finally, a future step is also to analyze how the proposed approach can be applied to data augmentation tasks in the context of event classification in videos.

REFERENCES

[1] Y. Kong and Y. Fu, "Human action recognition and prediction: A survey," *CoRR*, vol. abs/1806.11230, 2018. [Online]. Available: http://arxiv.org/abs/1806.11230

[2] A. F. Bobick and J. W. Davis, "The recognition of human movement using temporal templates," *IEEE Transactions on pattern analysis and machine intelligence*, vol. 23, no. 3, pp. 257–267, 2001.

[3] H. Venkateswara, S. Chakraborty, and S. Panchanathan, "Deep-learning systems for domain adaptation in computer vision: Learning transferable feature representations," *IEEE Signal Processing Magazine*, vol. 34, no. 6, pp. 117–129, 2017.

[4] S. Sun, H. Shi, and Y. Wu, "A survey of multi-source domain adaptation," *Information Fusion*, vol. 24, pp. 84–92, 2015.

[5] N. FarajiDavar, T. De Campos, J. Kittler, and F. Yan, "Transductive transfer learning for action recognition in tennis games," in *2011 IEEE International Conference on Computer Vision Workshops (ICCV Workshops)*. IEEE, 2011, pp. 1548–1553.

[6] L. Shao, F. Zhu, and X. Li, "Transfer learning for visual categorization: A survey," *IEEE transactions on neural networks and learning systems*, vol. 26, no. 5, pp. 1019–1034, 2014.

[7] A. Quattoni, M. Collins, and T. Darrell, "Transfer learning for image classification with sparse prototype representations," in *2008 IEEE Conference on Computer Vision and Pattern Recognition*. IEEE, 2008, pp. 1–8.

[8] L. Cao, Z. Liu, and T. S. Huang, "Cross-dataset action detection," in *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*. IEEE, 2010, pp. 1998–2005.

[9] T. J. Darrell, I. A. Essa, and A. P. Pentland, "Task-specific gesture analysis in real-time using interpolated views," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 18, no. 12, pp. 1236–1242, 1996.

[10] L. Wang, C. Gao, L. Yang, Y. Zhao, W. Zuo, and D. Meng, "Pm-gans: Discriminative representation learning for action recognition using partial-modalities," in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018, pp. 384–401.

[11] Y. Tas and P. Koniusz, "Cnn-based action recognition and supervised domain adaptation on 3d body skeletons via kernel feature maps," *arXiv preprint arXiv:1806.09078*, 2018.

[12] J. Gao, T. Zhang, and C. Xu, "I know the relationships: Zero-shot action recognition via two-stream graph convolutional networks and knowledge graphs," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 33, 2019, pp. 8303–8311.

[13] Y. Liu, Z. Lu, J. Li, T. Yang, and C. Yao, "Deep image-to-video adaptation and fusion networks for action recognition," *IEEE Transactions on Image Processing*, vol. 29, pp. 3168–3182, 2019.

[14] G. E. Hinton, S. Osindero, and Y.-W. Teh, "A fast learning algorithm for deep belief nets," *Neural Computation*, vol. 18, no. 7, pp. 1527–1554, 2006.

[15] H. Ali, S. N. Tran, E. Benetos, and A. S. d. Garcez, "Speaker recognition with hybrid features from a deep belief network," *Neural Computing and Applications*, vol. 29, no. 6, pp. 13–19, 2018.

[16] M. Roder, L. A. Passos, L. C. F. Ribeiro, B. C. Benato, A. X. Falcão, and J. P. Papa, "Intestinal parasites classification using deep belief networks," in *International Conference on Artificial Intelligence and Soft Computing*. Springer, 2020, pp. 242–251.

[17] G. E. Hinton, "Training products of experts by minimizing contrastive divergence," *Neural Computation*, vol. 14, no. 8, pp. 1771–1800, 2002.

[18] S. N. Gowda, "Human activity recognition using combinatorial deep belief networks," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, 2017, pp. 1–6.

[19] K. H. Ali and T. Wang, "Learning features for action recognition and identity with deep belief networks," in *2014 International Conference on Audio, Language and Image Processing*. IEEE, 2014, pp. 129–132.

[20] Y. Zhang, Z. Liu, and W. Zhou, "Event recognition based on deep belief network," *Acta Electronica Sinica*, vol. 45, no. 6, pp. 1415–1423, 2017.

[21] Y.-G. Jiang, S. Bhattacharya, S.-F. Chang, and M. Shah, "High-level event recognition in unconstrained videos," *International Journal of Multimedia Information Retrieval*, vol. 2, no. 2, pp. 73–101, 2013.

[22] J. Liu, B. Kuipers, and S. Savarese, "Recognizing human actions by attributes," in *Proceedings of the 2011 IEEE Conference on Computer Vision and Pattern Recognition*, ser. CVPR '11. IEEE Computer Society, 2011, pp. 3337–3344.

[23] C. Tan, F. Sun, T. Kong, W. Zhang, C. Yang, and C. Liu, "A survey on deep transfer learning," in *International conference on artificial neural networks*. Springer, 2018, pp. 270–279.

[24] K. Weiss, T. M. Khoshgoftaar, and D. Wang, "A survey of transfer learning," *Journal of Big data*, vol. 3, no. 1, pp. 1–40, 2016.

[25] P. Smolensky, "Parallel distributed processing: Explorations in the microstructure of cognition," D. E. Rumelhart, J. L. McClelland, and C. PDP Research Group, Eds. Cambridge, MA, USA: MIT Press, 1986, vol. 1, ch. Information Processing in Dynamical Systems: Foundations of Harmony Theory, pp. 194–281.

[26] K. Soomro, A. R. Zamir, and M. Shah, "Ucf101: A dataset of 101 human actions classes from videos in the wild," *arXiv preprint arXiv:1212.0402*, 2012.

[27] H. Kuehne, H. Jhuang, E. Garrote, T. Poggio, and T. Serre, "Hmdb: a large video database for human motion recognition," in *2011 International Conference on Computer Vision*. IEEE, 2011, pp. 2556–2563.

[28] J. Yue-Hei Ng, M. Hausknecht, S. Vijayanarasimhan, O. Vinyals, R. Monga, and G. Toderici, "Beyond short snippets: Deep networks for video classification," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2015, pp. 4694–4702.

[29] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.