

Transition Forests: Learning Discriminative Temporal Transitions for Action Recognition and Detection

Guillermo Garcia-Hernando
Imperial College London
ggarciah@imperial.ac.uk

Tae-Kyun Kim
Imperial College London
tk.kim@imperial.ac.uk

Abstract

A human action can be seen as transitions between one’s body poses over time, where the transition depicts a temporal relation between two poses. Recognizing actions thus involves learning a classifier sensitive to these pose transitions as well as to static poses. In this paper, we introduce a novel method called transitions forests, an ensemble of decision trees that both learn to discriminate static poses and transitions between pairs of two independent frames. During training, node splitting is driven by alternating two criteria: the standard classification objective that maximizes the discrimination power in individual frames, and the proposed one in pairwise frame transitions. Growing the trees tends to group frames that have similar associated transitions and share same action label incorporating temporal information that was not available otherwise. Unlike conventional decision trees where the best split in a node is determined independently of other nodes, the transition forests try to find the best split of nodes jointly (within a layer) for incorporating distant node transitions. When inferring the class label of a new frame, it is passed down the trees and the prediction is made based on previous frame predictions and the current one in an efficient and online manner. We apply our method on varied skeleton action recognition and online detection datasets showing its suitability over several baselines and state-of-the-art approaches.

1. Introduction

Recognizing and localizing human actions is an important and classic problem in computer vision [1, 7] with a wide range of applications including pervasive health-care, robotics, game control, etc. With recently introduced cost-effective depth sensors and reliable real-time body pose estimation [22], skeleton-based action recognition has become popular because of the advantage of pose features over raw RGB video approaches in both accuracy and efficiency [33].

Popular approaches for action recognition and localization include using generative models such as state-space models [14, 31]; or tackling it as a classification problem of either the whole sequence [26, 40], a small chunk of frames [10, 36] or deep recurrent models [9, 16]. The best performing methods focus either on modelling the temporal dynamics using time-series models [37] or recognizing key-poses [38], showing that both static and dynamic information are important cues for actions. Motivated by this, we consider decision forests [3], which have been widely adopted in computer vision [22, 24, 33], owing to many desired properties: clusters obtained in leaf nodes, scalability, robustness to overfitting, multiclass learning and efficiency.

The main challenge of using decision forests for temporal problems lies in dealing with temporal dependencies. Previous approaches encode the temporal variable in the feature space by stacking multiple frames [10], handcrafting temporal features [34, 40] or creating codebooks [34]. However, these methods require the temporal cues to be explicitly given instead of automatically learning them. Attempting to relieve this, [11, 33] add a temporal regression term and frames individually vote for an action center, breaking the temporal continuity and thus not fully capturing the temporal dynamics. [14] proposed a generative state-space without exploiting the benefit of having rich labelled data. [6] groups pairs of distant frames and grows trees using handcrafted split functions to cover different label transitions, with the difficulty of designing domain-specific functions and making the model complexity to increase with the number of labels.

In this work, we propose ‘transition forests’, an ensemble of randomized tree classifiers that learns both static pose information and temporal transitions in a discriminative way. Temporal dynamics are learned while training the forest (besides any temporal dependencies in the feature space) and predictions are made by taking into account previous predictions. Introducing previous predictions makes the learning problem more challenging as a consequence of the “chicken and egg” problem: making a decision in a node that depends on the decision in other nodes and vice

versa. To tackle this problem, we propose a training procedure that iteratively groups pairs of frames that have similar associated frame transitions and class label in a given level of the tree. We combine both static and transition information by randomly assigning nodes to be optimized by classification or transition criteria. In the end of tree growth, training frames arriving at leaf nodes represent effectively a class label and associated transitions. We found that adding such temporal relation in training helped to obtain more robust single frame predictions. Using single frames helped us in keeping the complexity low and being able to make online predictions, two crucial conditions to make our approach applicable to real life scenarios.

2. Related work

Skeleton-based action recognition. Generative models [14, 31, 32] such as Hidden Markov Models (HMM) have been proposed with the disadvantages of being difficult to estimate model parameters and time consuming learning and inference stages. Discriminative approaches have been widely adopted due to their superior performance and efficiency. For instance, [29] extracts local features from body joints captures temporal dynamics using Fourier Temporal Pyramids (FTP), further classifying the sequence using Support Vector Machines (SVM). Similarly, [26, 27] represents the whole skeletons as points in a Lie group before temporally aligning sequences using Dynamic Time Warping (DTW) and applying FTP. [36] proposes a Moving Pose descriptor (MP) using both pose and atomic motion information and then temporally mining key frames using a k-NN approach in contrast to [12] that uses DTW. Using key frames or key motion units has been also studied by [8, 28, 38] showing good performance revealing that static information is important to recognize actions. Recently, deep models using Recurrent Neural Networks (RNN) [9] and Long-Short Term Memory (LSTM) [25, 39] have been proposed to model temporal dependencies, but showed inferior performance than recent (offline) models that explicitly exploit static information [28, 30] or well-suited time-series mining [37]. Our forest learns best static per-frame and temporal information in a discriminative way.

Skeleton-based online action detection. Detecting actions on streaming data [7] has been less explored than recognizing segmented sequences, while being more interesting in real scenarios. Early approaches [10] include using short sequences of frames or short motion information [36] to vote if an action is being performed. A similar approach but adding multi-scale information was proposed by [20], while [17] proposed a dynamic bag of features. Recently, [16] introduced a more realistic dataset, baseline methods and shown state-of-the-art performance with a classification/regression RNN, later improved by [2] with the use of RGB-D spatio-temporal contexts and decision forests.

Forests and temporal data. Standard forest approaches for action recognition such as [10] directly stack frames and grows forests to classify them. [19, 40] create bags of poses and classified the whole sequences. Using the clustering properties of trees, [34] construct codebooks with the help of different heuristic rules capturing structural information. These approaches require the temporal cues to be directly encoded in the feature space. To relieve this, [4, 33, 35] add a temporal regression term and maps appearance and pose features to vote in an action Hough space. [11] proposes Trajectory Hough Forest (THF) that computes histograms of tree paths over consecutive color and flow trajectory patches and uses them as weights for prediction. However, in Hough frameworks, temporal information is captured as temporal offsets with respect to a temporal center of independent samples, breaking the temporal continuity and requiring the whole sequence to be observed. On the contrary, we explicitly capture the rich temporal dynamics and are able to perform online predictions. [6] proposes Pairwise Conditional Random Forests (PCRF) for facial expression recognition consisting of trees of which handcrafted split functions operate on pairs of frames. These pairs are formed to cover different facial dynamics and fed into multiple subsets of decision trees that are conditionally drawn based on different label transitions, making the ensemble size proportional to the number of labels. By contrast, our layer-wise optimization tries to automatically learn the best node splits based on single frames maximizing both static and transition information within the same tree and thus not needing handcrafted split functions or to create different trees based on different labels. Generative methods based on forests include Dynamic Forest Models (DFM) [14], which are ensembles of autoregressive trees that store multivariate distributions at their leaf nodes. These distributions model observation probabilities given short history of previous k frames. Similar to HMM, a decision forest is trained for each action label and inference is performed maximizing likelihood of the observed sequence. Recently, [5] proposed to learn smooth temporal regressors for real time camera planning. We share with [5] the recurrent nature of making online predictions conditioned on our own previous predictions, however our approach differs in how the recurrency is defined in both learning and inference stages. We compare some relevant methods in Section 4.

Tree-based methods for structured prediction. A related line of work [13, 18, 21, 23] proposes decision forests methods for image segmentation. The objective of these approaches is to obtain coherent pixel labels and, in order to connect multiple pixel predictions, decision forests are linked with probabilistic graphical models. While these methods focus on the spatial coherence of predictions in an image space, our method tries to capture discriminative changes of data/prediction in a temporal domain.

3. Transition forests

Suppose we are given a training set S composed of temporal sequences of input-output pairs $\{(x_1, y_1), \dots, (x_t, y_t)\}$ where x_t is a frame feature vector encoding pose information and y_t is its corresponding action label (or background in detection setting). Our objective is to infer y_t for every given x_t using decision trees. On a decision tree, an input instance x_t starts at the root and traverses different internal nodes until it reaches a leaf node. Each internal node $i \in \mathcal{N}$ contains a binary split function f with parameters θ_i deciding whether the instance should be directed to the left or to the right child nodes.

Consider the set of nodes $\mathcal{N}_l \subset \mathcal{N}$ at a level l of a decision tree. Let S_i denote the set of labeled training instances (x_t, y_t) that reached node i (see Fig. 1). For each pair of nodes $i, j \in \mathcal{N}_l$, we can compute the set of pairs of frames T_i^j that travel from node i to node j in d time steps as:

$$T_i^j = \{ \{ (x_{t-d}, y_{t-d}), (x_t, y_t) \} \mid (x_{t-d}, y_{t-d}) \in S_i \wedge (x_t, y_t) \in S_j \}, \quad (1)$$

where we term the set of pairs of frames T_i^j as *transitions from node i to j* . Note that T_i^j depends on frames that reached nodes i and j and time distance d . In order to capture different temporal patterns, we vary the distance d from one to a k -distant frame. In the following, we will refer to parameter k as the *temporal order of the transition forest*.

In the example shown in Fig. 1 we observe that the decision $f(\theta_0, S_0)$ is quite good as it separates S_0 in two sets, S_1 and S_2 , in which one action label predominates. If we examine the transitions associated to this split, we see that we obtain two pure sets, T_1^1 and T_2^2 , one mixed set T_2^1 and one empty set T_1^2 . Imagine now that we observe the ‘kick’ frame in S_1 and we would have to make a decision based on this split, we would certainly assign the wrong label ‘duck’ with an uncertainty of $2/3$. Alternatively, if we check the previous observed frame (in S_2) and inspect its associated transition T_2^1 , the uncertainty is now $1/2$ and thus we would be less inclined to make a wrong decision.

From the above example, we deduce that if we had obtained a better split and both child nodes were pure, we would certainly make a good decision by only looking at child nodes. However, good splits are difficult to learn if the temporal dynamics are not well captured on the feature space. On the other hand, if we had obtained a split that made transitions pure, we could also make a good decision. These observations motivate us to study how learning transitions between frames can help us to improve our predictions by introducing temporal information that was not available otherwise.

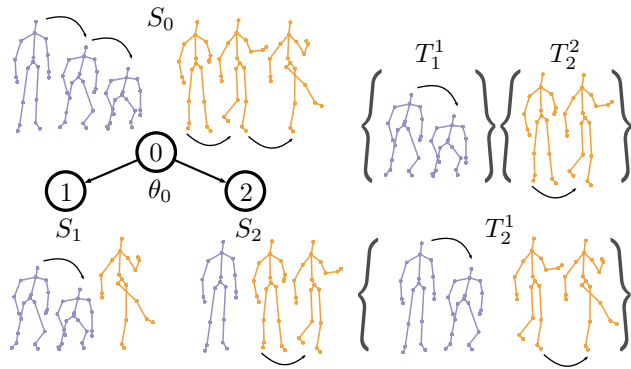


Figure 1: Consecutive frames representing two different actions (in purple ‘duck’, in orange ‘kick’) arrive at node 0. These frames are split in two different subsets S_1 and S_2 corresponding to child nodes 1 and 2. We compute the transitions as pairs of d -distant frames ($d = 1$ in this example) and we group them according to the route of each individual frame. T_1^1 and T_2^2 present only one transition, while T_2^1 two (one per class) and T_1^2 is empty. T_i^j are determined by θ_0 .

3.1. Learning transition forests

Our method for training a transition tree works by growing a tree one level at a time similar to [23]. At each level, we randomly assign one splitting criterion to each node, choosing between classification and transition. The classification criterion maximizes the class separation of static poses while the transition criterion groups frames that share similar transitions. As mentioned above, in order to maximize the span of temporal information learned, we learn transitions between d -distant pairs of frames (Eq. 1) from previous frame up to the temporal order of the forest, k . For each tree, we randomly assign a value of d in the mentioned range and we keep it constant during the growth of this particular tree. For a total ensemble of \mathcal{M} trees we will have subsets of trees trained with different d value: $\mathcal{M} = \mathcal{M}_1 \cup \dots \cup \mathcal{M}_k$.

Consider a node $i \in \mathcal{N}_l$ and a decision θ_i . According to θ_i , the instances in S_i are directed to its left or right child nodes, $2i+1$ and $2i+2$ respectively, as $S_{2i+1} = \{ (x_t, y_t) \in S_i \mid f(\theta_i, x_t) \leq 0 \}$ and $S_{2i+2} = S_i \setminus S_{2i+1}$. Note that the split function f operates on a single frame, which will be shown important in the inference stage. After splitting, we can compute the sets of transitions between their child nodes $\{2i+1, 2i+2\} \subseteq \mathcal{N}_{l+1}$ as $\{T_{2i+m}^{2i+n}\}_{m,n \in \{1,2\}}$. Note that T_i^i is split in four disjoint sets, each one related to the combination of transitions associated to its child nodes. The decision θ_i is chosen based on the minimization of an objective function.

Objective function. The objective function has two associated terms: one for single frame classification E_c and one for transitions between child nodes denoted as E_t . The

classification term E_c is the weighted Shannon entropy of the class distributions over the set of samples that reach the child nodes $\{S_{2i+m}\}_{m \in \{1,2\}}$ as in standard classification forests. Willing to decrease the uncertainty of transitions while growing the tree, the transition term aims to learn node decisions in a way that subsets of transitions are more pure in the next level. For a node $j \in \mathcal{N}_l$, the transition term is a function of the transitions between its child nodes and it is defined as:

$$E_t(\theta_j) = \sum_{m,n \in \{1,2\}} |T_{2j+m}^{2j+n}| H(T_{2j+m}^{2j+n}), \quad (2)$$

where $T_{(\cdot)}^{(\cdot)}$ is defined in Eq. (1) and $H(T_{(\cdot)}^{(\cdot)})$ is the Shannon entropy computed over the different label transitions. These two terms could be alternated or weighted-summed as single node optimizations. However, in order to reflect transitions between more distant nodes and capture further temporal information, we extend E_t to consider the set of all available nodes in a given level of a tree (as shown in Fig. 2 (a)). For this, we randomly assign a subset of parent nodes N_c and N_t to be optimized by E_c and E_t respectively. Given that transitions between nodes depend on the split decisions at different nodes, the task of learning a level can be formulated as the joint minimization of an objective function over the split parameters associated to the level nodes as:

$$\min_{\{\theta_i\}} E_c(\{\theta_i\}_{i \in N_c}) + E_t(\{\theta_i\}_{i \in N_c \cup N_t}). \quad (3)$$

Optimization. The problem of minimizing the objective function (Eq. 3) is hard to solve. One could think of randomly assign values to $\{\theta_i\}$ and pick the values that minimize the objective in a similar way to standard greedy optimization in decision trees. However, the search space grows exponentially with the depth of the tree and evaluating E_t for all nodes and samples at the same time is computationally expensive. Our strategy to relieve these problems is presented in Algorithm 1. Given that E_c only depends on decisions in N_c nodes, we can optimize these nodes using the standard greedy procedure. Once optimized and fixed all nodes in N_c , we iterate over every node in N_t to find the split function that minimizes a local version of E_t , denoted as E'_t , that keeps all the split parameters fixed except the one of the considered node. It is defined for a node $j \in N_t$ and it depends on the transitions between its child nodes and all the transitions *from* and *to* these child nodes:

$$\begin{aligned} E'_t(\theta_j | \{\theta_i\}_{i \neq j \in N_c \cup N_t}) &= \sum_{m,n \in \{1,2\}} \overbrace{|T_{2j+m}^{2j+n}| H(T_{2j+m}^{2j+n})}^{\text{between } j\text{'s child nodes (c.n.)}} \\ + \sum_{m,n \in \{1,2\}} &\underbrace{|T_{2j+m}^{2i+n}| H(T_{2j+m}^{2i+n})}_{\text{from } j\text{'s c.n. to } i\text{'s c.n.}} + \underbrace{|T_{2i+m}^{2j+n}| H(T_{2i+m}^{2j+n})}_{\text{to } j\text{'s c.n. from } i\text{'s c.n.}}. \end{aligned} \quad (4)$$

Algorithm 1 Learning level l of a transition tree

Input: Set of nodes \mathcal{N}_l at level l and temporal order d

Output: Set of split function parameters $\{\theta_i\}$

```

1: procedure LEARNLEVEL( $\mathcal{N}_l$ )
2:   randomly assign nodes in  $\mathcal{N}_l$  to  $N_c$  and  $N_t$ 
3:   for all  $i \in N_c$  do
4:     optimize  $N_c$  using  $E_c$ 
5:     save and fix  $\theta_i$ 
6:   end for
7:   initialize  $\{\theta_j\}$  for  $j \in N_t$ 
8:   while something changes do
9:     for all  $j \in N_t$  do
10:       $\Theta \leftarrow$  random feature/threshold selection
11:       $\theta_j \leftarrow \arg \min_{\theta' \in \Theta} E'_t(\theta_j | \{\theta_i\}_{i \neq j \in N_c \cup N_t})$ 
12:    end for
13:  end while
14: end procedure

```

The value of E'_t decreases (or does not change) at each iteration, thus indirectly minimizing E_t . Following this strategy it is not likely to reach a global minimum, but in practice we found that is effective to our problem. Note that computing Eq. 4 needs the split parameters in other nodes to be available, forcing us to initialize them before the first iteration. We found that an initialization of nodes using E_c helped the algorithm to converge faster than using a random initialization relieving us of computational cost.

3.2. Inference

Restricting ourselves to the set of leaf nodes \mathcal{L} , we assign each transition subset $\{T_i^j\}_{i,j \in \mathcal{L}}$ a conditional probability distribution over label transitions denoted $\pi_i^j(y_t | y_{t-d})$. This is different from classification forests where the *classification probability* $\pi_i(y_t)$ is estimated over all the set of training instances S_i that reached the leaf node i . Instead, we focus on subsets of transitions that depend on the leaf node (prediction) that previous d -distant frame reached. Note that the split function f is defined for a single frame, enabling us to perform individual frame predictions. For an ensemble of \mathcal{M}_d transition trees, we define a prediction function given two d -distant frames:

$$p_d(y_t | x_t, x_{t-d}, y_{t-d}) = \frac{1}{|\mathcal{M}_d|} \sum_{m \in \mathcal{M}_d} (\pi_{\ell(x_t)}^{\ell(x_{t-d})}(y_t | y_{t-d}))^{(m)}, \quad (5)$$

where $\ell(x_t)$ and $\ell(x_{t-d})$ are the leaf nodes reached by x_t and x_{t-d} at m -th tree respectively. We name this probability as *transition probability*. We combine the transition probability for different previous pairs of frames up to k with the *classification probability* (see Fig. 2 (b)). Combining the static classification probability with the temporal transition probability defines our final prediction equation

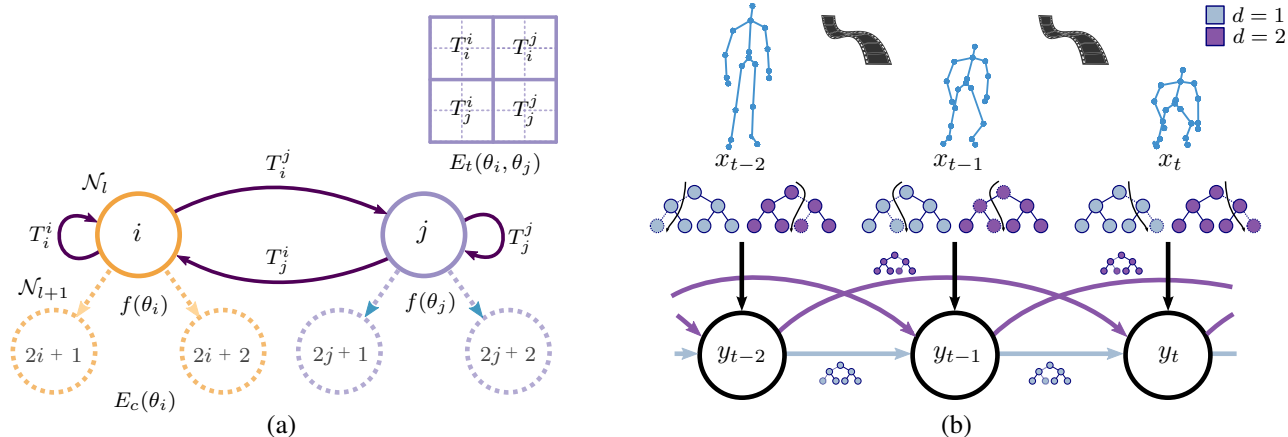


Figure 2: (a) Growing a level l of a transition tree depends on all the node decisions θ_i and θ_j at the same time. Each T_i^j divides in four disjoint sets according to the different routes that a pair of samples can follow. (b) In inference, each individual frame is passed down the forest and static pose classification is combined with transition probability. Transition probability is computed using the trees trained for specific d -distant frames (shown in different color). In this example $k = 2$ and $|\mathcal{M}| = 2$.

for a transition forest of temporal order k :

$$p(y_t | x_t, x_{t-1}, \dots, x_{t-k}, y_{t-1}, \dots, y_{t-k}) = \frac{1}{|\mathcal{M}|} \sum_m (\pi_{\ell(x_t)}(y_t))^{(m)} \frac{1}{k} \sum_{1 \leq d \leq k} p_d(y_t | x_t, x_{t-d}, y_{t-d}). \quad (6)$$

For each frame x_t we obtain a probability of the frame belonging to one action (plus background in detection setting) based on k previous predictions. In the *action recognition* setting we average the per-frame results to predict the whole sequence. On the other hand, for *online action detection*, we define two thresholds, β_s and β_e , to locate the start and the end frame of the action. When the score for one action exceeds β_s , we aggregate the results since the start of the action and we do not allow any action change until the score is less than β_e .

3.3. Implementation details

If the training data is not enough, we may encounter empty transition subsets at low levels of the tree. For this reason, we set a minimum number of instances needed to estimate their probability distribution and we empirically set this parameter to ten in our experiments. This parameter is conceptually the same as the stopping criterion of requiring a minimum number of samples to keep splitting a node.

4. Experimental evaluation

In the following we present experiments to evaluate the effectiveness of our approach. We start evaluating our approach for action recognition and we follow with online action detection. In all experiments we performed standard

pre-processing on given joint positions similar to [26] making them invariant to scale, rotation and point of view.

4.1. Baselines

We compare our approach with five different forest-based baselines detailed next. For fair comparison, we always use the same number of trees in all methods and we adjust the maximum depth for best performance.

Random Forest [3] (RF). To assess how well performs a decision forest while only using static information, we implement a single frame-based random forest only using E_c .

Sliding Window Forest [10] (SW). To compare our learning of temporal dynamics with the strategy of stacking multiple frames, we implement a forest using the sliding window setting in which the temporal order k the number of previous frames in the window.

Trajectory Hough Forest [11] (THF). To compare with a temporal regression method, we implement [11] and adapt their color trajectories to poses and their histograms to deal with a temporal order of k .

Dynamic Forest Model [14] (DFM). To compare our discriminative forest approach with a generative one, our third baseline is the a generative forest where k is the order of their non-linear Markov model. With no public implementation available, we directly report results in [14].

Pairwise Conditional Random Forest [6] (PCRF). To assess the discriminative pairwise information, we implement a pairwise forest similar to the one used for expression recognition [6]. We grow and combine classification trees for different pairwise temporal distance up to k .

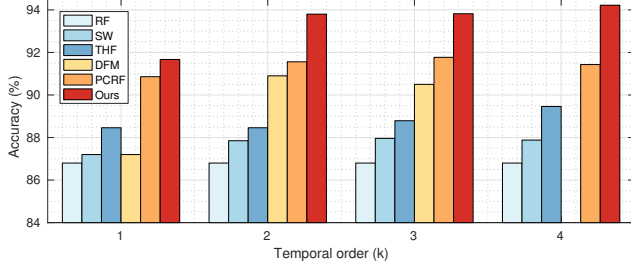


Figure 3: Temporal order k for different baselines and our approach on MSRC-12 dataset.

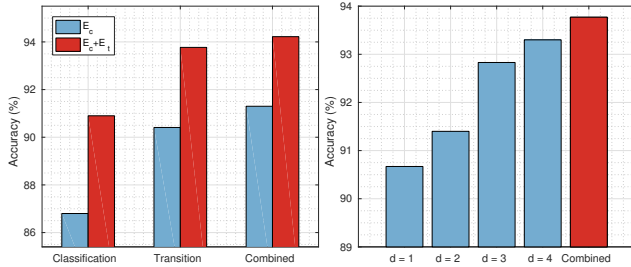


Figure 4: (a) E_c vs. $E_c + E_t$ and terms in Eq. 6. (b) contribution of different d order trees to transition probability shown in (a) and defined in Eq. 5 on MSRC-12.

4.2. Action recognition experiments

We evaluate the proposed algorithm on three different action recognition benchmarks: MSRC-12 [10], MSR-Action3D [15] and Florence-3D [19]. First, we perform detailed control experiments and parameter evaluation on MSRC-12 dataset. Next, we evaluate our approach comparing with baselines and state-of-the-art on all datasets.

4.2.1 MSRC-12 experiments

The MSRC-12 [10] dataset consists of 12 iconic and metaphoric gestures performed by 30 different actors. We follow the experimental protocol in [14]: only the 6 iconic gestures are used, making a total of 296 sequences and we perform 5-fold leave-person-out cross-validation, *i.e.*, 24 actors for training and 6 actors for testing per fold.

Temporal order k and comparison with baselines. In Fig. 3 we show experimental results varying the temporal order parameter k for all approaches. We observe that using only static information on single frames (RF) to recognize action is limited and it can be improved by stacking multiple frames (SW). Adding a regression term as in THF helps to increase the accuracy. DFM uses the same exact input window as SW, while being more robust as a result of their explicit modeling of time. Better than the rest of baselines, PCRF shows that capturing pairwise information is effective to model the temporal dynamics of the actions. On the

Method	Year	Real-time	Online	Acc (%)
DFM [14]	2014	✓	✓	90.90
ESM [12]	2014	✗	✗	96.76
Riemann [8]	2015	✗	✗	91.50
PCRF (our result) [6]	2015	✓	✓	91.77
Bag-of-poses [38]	2016	✗	✗	94.04
Ours (JP)	2016	✓	✓	94.22
Ours (RJP)	2016	✓	✓	97.54
Ours (MP)	2016	✓	✓	98.25

Table 1: MSRC-12: Comparison with state-of-the-art using different frame representations.

other hand, our approach shows the best performance for all temporal orders. This shows that both combining static and temporal information in a discriminative way is very effective. In the next two paragraphs we analyze the contribution of both sources of information.

Discriminative power of learned transitions. We measure the impact of our transition training procedure presented in Section 3.1. For this, we train two different transition forests, one using only E_c and one using E_c and E_t . For each forest, we show the performance by breaking down the terms of Eq. 6: (i) using only the classification probability; (ii) using only the transition probability (Eq. 5); (iii) combining both terms (Eq. 6).

Results are shown in Fig. 4 (a). We observe that our proposed training algorithm increases the performance of both static and transition terms, leading to an important overall improvement. The static classification term improves substantially, meaning that E_t helps to separate categories on the feature space by introducing temporal information that was not available otherwise. In Fig. 4 (b) we show the contribution of each temporal distance to the overall transition probability in Eq. 5.

Frame representation. In addition to joint positions (JP) from above experiments, we experimented with two different frame representations: one static and one dynamic. The static one consists of pairwise relative distance of joints (RJP), proven to be more robust than JP while being very simple [26]. The dynamic one, named Moving Pose (MP) [36] incorporates temporal information by adding velocity and acceleration of joints using nearby frames. In Table 1 we observe that RJP and MP perform similarly well performing better than JP, showing that our approach can benefit of different static and dynamic feature representations.

Initialization. We initialized the transition nodes N_t in two ways: randomly and using E_c . We found that the latter initialization provided slightly better results by 0.35% after ten iterations. However, after doubling the number of iterations, the difference was reduced to 0.07%, leading to the conclusion that our algorithm is robust to initialization, but correctly initializing reduces the training time. Based

on this, we limited the number of iterations to ten.

Ensemble size. A single tree of maximum depth 10 gave us an accuracy of 86.42%, six trees 93.10% and twelve 94.22%. As a tree-based algorithm, adding more trees is expected to increase the performance (up to saturation) at the cost of computational time.

Comparison with the state-of-the-art. In Table 1 we compare our approach with the state-of-the-art. We observe that using the simple JP representation, we achieve the best with the exception of ESM [12]. However, ESM uses a slow variant of DTW and MP representation. Using both RJP and MP representation our approach achieves the best performance while being able to run in real time (1778 fps).

4.2.2 MSR-Action3D experiments.

The MSR-Action3D [15] dataset is composed of 20 actions performed by 10 different actors. Each actor performed every action two or three times for a total of 557 sequences. We perform our main experiments following the setting proposed by [15]. In this protocol, the dataset is divided into three subsets of eight actions, named AS1, AS2 and AS3. The classification is performed on each subset separately and the final classification accuracy is the average over the three subsets. We perform a cross-subject validation in which half of the actors are used for training and the rest for testing using ten different splits. We use RJP frame representation, $k = 4$ and 50 trees of maximum depth 8.

Baselines and state-of-the-art comparison are shown in Tables 2 and 3 respectively. Our approach achieves better performance than all baselines. Offline state-of-the-art methods [28, 37] achieve the best performance. Focusing on methods that are both real-time and online, the best performance is achieved by HURNN-L [9], which uses a deep architecture to learn an end-to-end classifier. We obtain better results than [9] on both their online and offline flavors.

Some authors [25, 36] show results using a different protocol [29] in which all 20 actions are considered. For comparison, using this protocol we achieved an accuracy of 92.8%, which is superior to state-of-the-art online approaches of MP [36], 91.7%, and dLSTM [25], 92.0%, but inferior to the offline approach of Gram matrix [37], 94.7%. It is important to note that the inference complexity of both [36, 37] increases with the number of different actions, which is not the case of our approach, making it more suitable for realistic scenarios. [37] reported a testing time (ten runs over whole testing set) of 1523 seconds, for the same setting we report a significant lower time of 289 s.

4.2.3 Florence-3D experiments

The Florence-3D dataset [19] consists of 9 different actions performed by 10 subjects. Each subject performed every

Method	MSRC-12	MSR-Action3D	Florence-3D
RF [3]	86.83	87.77	85.46
SW [10]	87.81	90.48	88.44
THF [11]	89.46	91.31	89.06
DFM [14]	90.90	-	-
PCRF [6]	91.77	92.09	91.23
Ours	94.22	94.57	94.16

Table 2: Comparison with forest-based baselines.

Method	Year	Real-time	Online	Acc (%)
Bag of poses [19]	2013	✗	✗	82.15
Lie group [26]	2014	✗	✗	90.88
PCRF (our result) [6]	2015	✓	✓	91.23
Rolling rot. [27]	2016	✗	✗	91.40
Graph-based [30]	2016	✗	✗	91.63
Key-poses [28]	2016	✓	✗	92.25
Ours	2016	✓	✓	94.16

Table 4: Florence-3D: Comparison with state-of-the-art.

action two or three times making a total of 215 action sequences. Following previous work [28, 30], we adopt a leave-one-subject-out protocol, *e.g.* nine subjects are used for training and one for testing for ten times. We used the same parameters as in the previous experiment.

We compare the proposed approach with baselines and state-of-the-art in Tables 2 and 4 respectively. We can see that our approach achieves the best performance over all baselines and state-of-the-art. Note that on this dataset we outperform the recent Key-poses approach [28], which achieved the best performance on MSR-Action3D dataset.

4.3. Online action detection experiments

We end our experimental evaluation in a more realistic scenario. We test our approach for online action detection on the very recently proposed Online Action Detection (OAD) dataset [16]. The dataset consists of 59 long sequences containing 10 different daily-life actions performed by different actors. Each sequence contains different action/background periods of variable length in arbitrary order annotated with start/end frames. We use the same splits and evaluation protocol as [16]. Previous work [16] fixed the number of considered previous frames to 10, in consequence we set $k = 10$. We use RJP representation and 50 trees of maximum depth 20. Thresholds β_s and β_e were empirically set to 0.79 and 0.16 respectively.

In Table 5 we report class-wise and overall F1-score for baselines, state-of-the-art and our approach. We also report the accuracy of start and end frame detection ‘SL’ and ‘EL’ respectively. We observe that our approach outperforms all baselines. PCRF forest shown the best results among the

Method	Year	Real-time	Online	AS1 (%)	AS2 (%)	AS3 (%)	Average (%)
BoF forest [40]	2013	✗	✗	-	-	-	90.90
Lie group [26]	2014	✗	✗	95.29	83.87	98.22	92.46
HBRNN-L [9]	2015	✓	✗	93.33	94.64	95.50	94.49
Graph-based [30]	2016	✗	✗	93.75	95.45	95.10	94.77
Gram matrix [37]	2016	✓	✗	98.66	94.11	98.13	96.97
Key-poses [28]	2016	✓	✗	-	-	-	97.44
PCRF (our result) [6]	2015	✓	✓	94.51	85.58	96.18	92.09
HURNN-L [9]	2015	✓	✓	92.38	93.75	94.59	93.57
Ours	2016	✓	✓	96.10	90.54	97.06	94.57

Table 3: MSR-Action3D: Comparison with state-of-the-art.

Action	Baselines			State-of-the-art		
	RF	SW	PCRF	RNN [39]	JCR-RNN [16]	Ours
drinking	0.598	0.387	0.468	0.441	0.574	0.705
eating	0.683	0.590	0.550	0.550	0.523	0.700
writing	0.640	0.678	0.703	0.859	0.822	0.758
opening cupboard	0.367	0.317	0.303	0.321	0.495	0.473
washing hands	0.698	0.792	0.613	0.668	0.718	0.740
opening microwave	0.525	0.717	0.717	0.665	0.703	0.717
sweeping	0.539	0.583	0.635	0.590	0.643	0.645
gargling	0.298	0.414	0.464	0.550	0.623	0.633
throwing trash	0.340	0.205	0.350	0.674	0.459	0.518
wiping	0.823	0.765	0.823	0.747	0.780	0.823
Overall	0.578	0.556	0.607	0.600	0.653	0.712
SL	0.361	0.366	0.378	0.366	0.418	0.514
EL	0.391	0.326	0.412	0.376	0.443	0.527
Inference time (s)	0.59	0.61	3.58	3.14	2.60	1.84

Table 5: Performance comparison on Online Action Detection (OAD) dataset.

baselines with a performance comparable to RNN, showing that temporal pairwise information is important. On the other hand, RF performs particularly well on this dataset, revealing that distinguishing static poses is important in addition to temporal information. Combining both static and temporal information in our approach led us to better performance than the current state-of-the-art JCR-RNN [16], which added a regression term on a LSTM to predict both start and end frames of actions.

Efficiency. We measure the average inference time on 9 long sequences of 3200 frames in average. We present the results at the bottom of Table 5 with a C++ implementation on a Intel Core i7 (2.6 GHz) and 16 GB RAM. All compared approaches are real-time, with JCR-RNN achieving 1230 fps for 1778 fps of our approach, showing that we can obtain high performance while keeping the complexity low.

5. Summary and conclusion

We proposed a new forest based classifier that is able to learn both static poses and transitions in a discriminative way. Our proposed training procedure helps to capture temporal dynamics in a more effective way than other strong forest baselines. Introducing temporal relationships while growing the trees and also using them in inference helped to obtain more robust frame-wise predictions, leading us to show state-of-the-art performance in both challenging problems of action recognition and online action detection.

Currently, our learning stage is limited to pairwise transitions and we believe that it would be interesting to incorporate different time orders within the same tree learning. Also, given the generality of our work, it would be interesting to test its performance using other data modalities (such as RGB/depth frame features) or applied to other temporal problems requiring efficient and online classification.

References

- [1] J. K. Aggarwal and M. S. Ryoo. Human activity analysis: A review. *ACM Computing Surveys*, 2011. [1](#)
- [2] S. Baek, K. I. Kim, and T.-K. Kim. Real-time online action detection forests using spatio-temporal contexts. In *WACV*, 2017. [2](#)
- [3] L. Breiman. Random forests. In *Machine learning*, 2001. [1](#), [5](#), [7](#)
- [4] H. J. Chang, G. Garcia-Hernando, D. Tang, and T.-K. Kim. Spatio-temporal hough forest for efficient detection-localisation-recognition of fingerwriting in egocentric camera. In *CVIU*, 2016. [2](#)
- [5] J. Chen, H. M. Le, P. Carr, Y. Yue, and J. J. Little. Learning online smooth predictors for realtime camera planning using recurrent decision trees. In *CVPR*, 2016. [2](#)
- [6] A. Dapogny, K. Bailly, and S. Dubuisson. Pairwise conditional random forests for facial expression recognition. In *ICCV*, 2015. [1](#), [2](#), [5](#), [6](#), [7](#), [8](#)
- [7] R. De Geest, E. Gavves, A. Ghodrati, Z. Li, C. Snoek, and T. Tuytelaars. Online action detection. In *ECCV*, 2016. [1](#), [2](#)
- [8] M. Devanne, H. Wannous, P. Pala, S. Berretti, M. Daoudi, and A. Del Bimbo. Combined shape analysis of human poses and motion units for action segmentation and recognition. In *FG*, 2015. [2](#), [6](#)
- [9] Y. Du, W. Wang, and L. Wang. Hierarchical recurrent neural network for skeleton based action recognition. In *CVPR*, 2015. [1](#), [2](#), [7](#), [8](#)
- [10] S. Fothergill, H. Mentis, P. Kohli, and S. Nowozin. Instructing people for training gestural interactive systems. In *ACM SIGCHI*, 2012. [1](#), [2](#), [5](#), [6](#), [7](#)
- [11] G. Garcia-Hernando, H. J. Chang, I. Serrano, O. Deniz, and T.-K. Kim. Transition hough forest for trajectory-based action recognition. In *WACV*, 2016. [1](#), [2](#), [5](#), [7](#)
- [12] H.-J. Jung and K.-S. Hong. Enhanced sequence matching for action recognition from 3d skeletal data. In *ACCV*. 2014. [2](#), [6](#), [7](#)
- [13] P. Kotschieder, P. Kohli, J. Shotton, and A. Criminisi. Geof: Geodesic forests for learning coupled predictors. In *CVPR*, 2013. [2](#)
- [14] A. M. Lehrmann, P. V. Gehler, and S. Nowozin. Efficient nonlinear markov models for human motion. In *CVPR*, 2014. [1](#), [2](#), [5](#), [6](#), [7](#)
- [15] W. Li, Z. Zhang, and Z. Liu. Action recognition based on a bag of 3d points. In *CVPRW*, 2010. [6](#), [7](#)
- [16] Y. Li, C. Lan, J. Xing, W. Zeng, C. Yuan, and J. Liu. Online human action detection using joint classification-regression recurrent neural networks. In *ECCV*, 2016. [1](#), [2](#), [7](#), [8](#)
- [17] M. Meshry, M. E. Hussein, and M. Torki. Linear-time online action detection from 3d skeletal data using bags of gesturelets. In *WACV*, 2016. [2](#)
- [18] S. Nowozin, C. Rother, S. Bagon, T. Sharp, B. Yao, and P. Kohli. Decision tree fields. In *ICCV*, 2011. [2](#)
- [19] L. Seidenari, V. Varano, S. Berretti, A. Bimbo, and P. Pala. Recognizing actions from depth cameras as weakly aligned multi-part bag-of-poses. In *CVPRW*, 2013. [2](#), [6](#), [7](#)
- [20] A. Sharaf, M. Torki, M. E. Hussein, and M. El-Saban. Real-time multi-scale action detection from 3d skeleton data. In *WACV*, 2015. [2](#)
- [21] J. Shotton, M. Johnson, and R. Cipolla. Semantic texton forests for image categorization and segmentation. In *CVPR*, 2008. [2](#)
- [22] J. Shotton, T. Sharp, A. Kipman, A. Fitzgibbon, M. Finocchio, A. Blake, M. Cook, and R. Moore. Real-time human pose recognition in parts from single depth images. *Communications of the ACM*, 2013. [1](#)
- [23] J. Shotton, T. Sharp, P. Kohli, S. Nowozin, J. Winn, and A. Criminisi. Decision jungles: Compact and rich models for classification. In *NIPS*, 2013. [2](#), [3](#)
- [24] D. Tang, T.-H. Yu, and T.-K. Kim. Real-time articulated hand pose estimation using semi-supervised transductive regression forests. In *ICCV*, 2013. [1](#)
- [25] V. Veeriah, N. Zhuang, and G.-J. Qi. Differential recurrent neural networks for action recognition. In *ICCV*, 2015. [2](#), [7](#)
- [26] R. Vemulapalli, F. Arrate, and R. Chellappa. Human action recognition by representing 3d skeletons as points in a lie group. In *CVPR*, 2014. [1](#), [2](#), [5](#), [6](#), [7](#), [8](#)
- [27] R. Vemulapalli and R. Chellappa. Rolling rotations for recognizing human actions from 3d skeletal data. In *CVPR*, 2016. [2](#), [7](#)
- [28] C. Wang, Y. Wang, and A. L. Yuille. Mining 3d key-pose-motifs for action recognition. In *CVPR*, 2016. [2](#), [7](#), [8](#)
- [29] J. Wang, Z. Liu, Y. Wu, and J. Yuan. Mining actionlet ensemble for action recognition with depth cameras. In *CVPR*, 2012. [2](#), [7](#)
- [30] P. Wang, C. Yuan, W. Hu, B. Li, and Y. Zhang. Graph based skeleton motion representation and similarity measurement for action recognition. In *ECCV*, 2016. [2](#), [7](#), [8](#)
- [31] D. Wu and L. Shao. Leveraging hierarchical parametric networks for skeletal joints based action segmentation and recognition. In *CVPR*, 2014. [1](#), [2](#)
- [32] L. Xia, C.-C. Chen, and J. Aggarwal. View invariant human action recognition using histograms of 3d joints. In *CVPRW*, 2012. [2](#)
- [33] A. Yao, J. Gall, G. Fanelli, and L. J. Van Gool. Does human action recognition benefit from pose estimation?. In *BMVC*, 2011. [1](#), [2](#)
- [34] T.-H. Yu, T.-K. Kim, and R. Cipolla. Real-time action recognition by spatiotemporal semantic and structural forests. In *BMVC*, 2010. [1](#), [2](#)
- [35] T.-H. Yu, T.-K. Kim, and R. Cipolla. Unconstrained monocular 3d human pose estimation by action detection and cross-modality regression forest. In *CVPR*, 2013. [2](#)
- [36] M. Zanfir, M. Leordeanu, and C. Sminchisescu. The moving pose: An efficient 3d kinematics descriptor for low-latency action recognition and detection. In *ICCV*, 2013. [1](#), [2](#), [6](#), [7](#)
- [37] X. Zhang, Y. Wang, M. Gou, M. Sznai, and O. Camps. Efficient temporal sequence comparison and classification using gram matrix embeddings on a riemannian manifold. In *CVPR*, 2016. [1](#), [2](#), [7](#), [8](#)
- [38] G. Zhu, L. Zhang, P. Shen, and J. Song. Human action recognition using multi-layer codebooks of key poses and atomic motions. In *Signal Proc.: Image Comm.*, 2016. [1](#), [2](#), [6](#)
- [39] W. Zhu, C. Lan, J. Xing, W. Zeng, Y. Li, L. Shen, and X. Xie. Co-occurrence feature learning for skeleton based action recognition using regularized deep lstm networks. In *AAAI*, 2016. [2](#), [8](#)
- [40] Y. Zhu, W. Chen, and G. Guo. Fusing spatiotemporal features and joints for 3d action recognition. In *CVPRW*, 2013. [1](#), [2](#), [8](#)