
Temporal Gaussian Mixture Layer for Videos

AJ Piergiovanni¹ Michael S. Ryoo¹

Abstract

We introduce a new convolutional layer named the *Temporal Gaussian Mixture* (TGM) layer and present how it can be used to efficiently capture longer-term temporal information in continuous activity videos. The TGM layer is a temporal convolutional layer governed by a much smaller set of parameters (e.g., location/variance of Gaussians) that are fully differentiable. We present our fully convolutional video models with multiple TGM layers for activity detection. The extensive experiments on multiple datasets, including Charades and MultiTHUMOS, confirm the effectiveness of TGM layers, significantly outperforming the state-of-the-arts¹.

1. Introduction

Activity videos are spatio-temporal data: they are image frames with a specific width/height (XY) concatenated along time axis (T). Recognition from such videos requires capturing both spatial and temporal information, desirably using learned convolutional kernels. Temporal convolution is particularly beneficial in activity ‘detection’ tasks, which require making activity decisions at every frame given a *continuous* video (Sigurdsson et al., 2016b; Yeung et al., 2015). Previous methods investigated using 3-D XYT convolutional filters (Tran et al., 2014; Carreira & Zisserman, 2017) as well as the models with 2-D XY conv. layers followed by 1-D temporal conv. (Tran et al., 2018), pooling or attention layers (Piergiovanni et al., 2017).

Understanding complex multi-activity videos requires capturing information in long-term time intervals. Different frames contain different information, and the model needs to learn to take advantage of as many frames as possible, while abstracting them efficiently. Previous attempts of simply pooling representations over time or learning temporal

¹Department of Computer Science, Indiana University. Correspondence to: AJ Piergiovanni <ajpiergi@indiana.edu>, Michael Ryoo <mryoo@indiana.edu>.




Figure 1. Example illustrating how our Temporal Gaussian Mixture layer is computed. Multiple (M) temporal Gaussian distributions are learned, and they are combined with the learned soft attention (mixing) weights to form the C temporal convolution filters. L is the temporal length of the filter.

conv. filters with a small number of frames (e.g., 16 or 64) was thus often insufficient to fully consider rich long-term temporal context. Simultaneously, bruteforcely increasing the temporal filter length (to look at more frames) results more learnable parameters, requiring more training data, which can be expensive when activities are rare.

In this paper, we introduce a new convolutional layer named the *Temporal Gaussian Mixture* (TGM) layer, and present how it can be used to efficiently capture longer-term temporal information in activity videos. Our temporal Gaussian mixture layer is a temporal convolutional layer, whose filters/kernels are controlled by a set of (temporal) Gaussian distribution parameters. Each of our temporal Gaussian distributions specify (temporally) ‘where’ the model should look, and our Gaussian mixture layer combines them as multiple convolutional filters to be applied on top of temporally-continuous representations. This layer allows the video representation at each time step to be constructed while focusing on different neighboring temporal regions, instead of only focusing on its local segment. It is a convolutional layer governed by a much smaller set of parameters (i.e., locations/variances of the Gaussians as well as their mixture weights) that are fully differentiable. Importantly, the number of parameters of the TGM layer is independent of the length of the filter, allowing the model to capture longer-term temporal information without additional parameters.

The motivation behind our temporal Gaussian mixture layer is to learn the temporal structure of an activity as a composition of temporal Gaussian regions/attentions. Such structure allows the model to obtain a compact spatio-temporal representation abstracting each (long-term) time interval, using multiple temporal conv. layers with far fewer parameters. It is also related to the previous temporal attention works

(Piergiovanni et al., 2017), but our model is designed to be fully convolutional to handle continuous data and it learns more compositional structures with multiple layers.

We present video-CNN models using our TGM layers for activity detection in continuous videos. Our model stacks TGM layers on top of several state-of-the-art CNNs such as I3D (Carreira & Zisserman, 2017). This enables our model to capture longer-term temporal information than provided by the base CNNs, compositionally modeling the temporal structure with multiple TGM layers. Our model was evaluated on multiple public datasets including MultiTHUMOS and Charades, and was able to outperform the best previous activity detection CNNs by a meaningful margin.

2. Related Works

Learning video representations for human activity recognition has been successful. CNN methods allow end-to-end learning of video features and representations optimized for the training data, performing superior to traditional works (Aggarwal & Ryoo, 2011) for video understanding.

Two-stream CNN models take a single RGB frame and a small number of optical flow frames as inputs to capture both motion and appearance information in videos (Simonyan & Zisserman, 2014; Feichtenhofer et al., 2016). Models learning 3-D spatio-temporal (XYT) convolutional filters were designed and applied to many activity recognition tasks as well (Tran et al., 2014; Carreira & Zisserman, 2017; Tran et al., 2017; Hara et al., 2017). Large scale datasets for activity detection, such as THUMOS (Jiang et al., 2014), ActivityNet (Heilbron et al., 2015), Kinetics (Kay et al., 2017), and Charades (Sigurdsson et al., 2016b) provided these approach the necessary training data to learn the models. Such 3-D XYT CNNs were also used to capture spatio-temporal information for activity detection (Xu et al., 2017; Shou et al., 2016; 2017; Zhao et al., 2017). However, all these CNNs were limited to the consideration of a fixed local video segment (e.g., 16 frames in (Tran et al., 2014) and 64-99 frames in (Carreira & Zisserman, 2017)) when making activity decisions.

Some works studied combining representations over longer-term temporal intervals (Karpathy et al., 2014; Ng et al., 2015; Varol et al., 2017), but it was generally done with a temporal pooling of local representations or spatio-temporal convolutions with slightly larger fixed intervals. Recurrent neural networks (RNNs) have also been used to model activity transitions between frames (Yeung et al., 2015; 2016; Escorcia et al., 2016), but they were strictly sequential and had limitations in maintaining temporal information over a longer temporal duration, particularly for videos with multiple complex activities. Recently, CNN models using temporal attention for activity videos (Piergiovanni et al., 2017; Piergiovanni & Ryoo, 2018b) were studied as well.

However, a fully convolutional model to analyze continuous videos while efficiently representing information in long term intervals has been lacking.

Several works have explored using parameterized convolutional kernels for point clouds (Xu et al., 2018) or images (Cohen et al., 2018), but were quite different from our design to handle long-term temporal information.

Our layer is different from the previous standard spatio-temporal convolutional layers in that it relies on significantly fewer parameters by forcing filter shapes to be Gaussian compositions. Our temporal layer is also different from previous Gaussian Mixture Model layers (Variani et al., 2015) in that our layer is convolutional while they are not.

3. Approach

In this section, we introduce a new convolutional layer named the *Temporal Gaussian Mixture* (TGM) layer, and present how it can be used for activity recognition. Our Temporal Gaussian Mixture layer is a temporal convolutional layer to be applied on top of a sequence of representations (usually from frame-level or segment-level CNNs), whose filters/kernels are controlled by a set of (temporal) Gaussian distribution parameters. The motivation is to make each temporal Gaussian distribution specify (temporally) ‘where to look’ with respect to the activity center, and represent the activity as a collection/mixture of such temporal Gaussians convolved with video features. Our layer is fully differentiable and trainable using standard backpropagation.

Our TGM layer can be interpreted as a form of 1-D convolution where the filters are determined by a mixture of Gaussians. However, we design our TGM layer differs from the standard temporal convolutional layers of learning 1-D (time) or 2-D (channel-by-time) filters in the following aspects (illustrated in Fig. 2):

1. Instead of learning temporal convolution filters of any arbitrary values, our filter is forced to have the form of a temporal Gaussian mixture shared across all frame-level channels. This allows the layer to rely on significantly fewer number of (fully differentiable) parameters, while capturing the concept of temporal structure/attention.
2. Our temporal Gaussian mixture layer handles multiple 3-D tensors internally to preserve channels from the frame-level CNN by adding a new temporal channel axis. Its input is 3-D (channel-by-channel-by-time), where one channel dimension is inherited from the frame-level CNN and this dimension size remains unchanged. This is a form of grouped convolution, and we call this more specifically temporal channel grouping (TC-grouping).

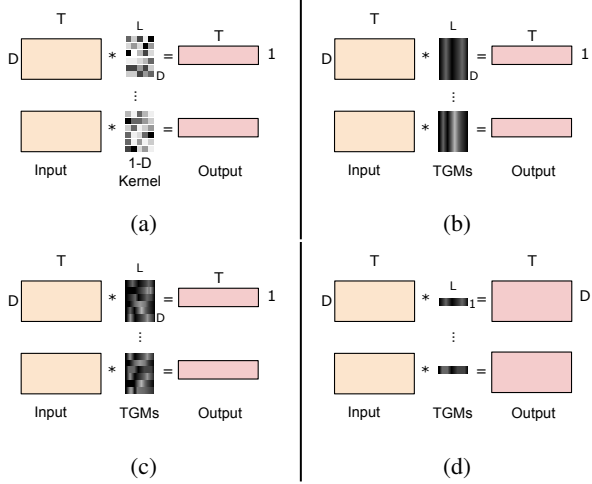


Figure 2. (a-c) Different forms of 1-D temporal convolutions which take a $D \times T$ input and produces a $C \times T$ output based on C number of $D \times L$ kernels: (a) the standard 1-D convolution, (b) using Gaussian mixtures for 1-D convolution while sharing Gaussian mixtures across input channels, and (c) using D different Gaussian mixtures for 1-D convolution. (d) Our TGM layer (with TC-grouping) in its simplest form (i.e., 1-layer case) applying the $1 \times L$ temporal kernel in a 2-D convolutional fashion, maintaining both time and feature axis.

3.1. Temporal Gaussian Mixture layer

Our temporal Gaussian mixture layer takes a 3-D input with the dimensionality of $C_{in} \times D \times T$, where C_{in} is the number of input channels, D is the dimensionality of the representations from frame-level (or segment-level) CNNs, and T is the time. Given such input, the TGM layer convolves it with C_{out} number of $1 \times L$ filters/kernels, generating a $C_{out} \times D \times T$ -dim representation as an output. L is the temporal length of the temporal Gaussian mixture filter. D is usually 1K or 4K and T is the number of time steps (frames) in each video (i.e., it varies per video). C_{out} is the number of different mixtures, corresponding to the number of output channels in standard convolution. In Fig. 2, we compare our TGM layer with the different forms of temporal convolutions. We experimentally compare all these versions to confirm the benefits of our TGM layer.

Our layer is composed of a set of M Gaussians. Each Gaussian has 2 parameters: a center $\hat{\mu}$ and a width $\hat{\sigma}$. Each layer has additional hyper-parameters: L , the temporal duration and M , the number of Gaussians to learn. We constrain the learned center to be between 0 and L and σ to be positive:

$$\mu = (L - 1) \cdot \frac{\tanh(\hat{\mu}) + 1}{2}, \quad \sigma^2 = \exp(\hat{\sigma}). \quad (1)$$

We use the above μ and σ to construct the temporal Gaussian kernels. This acts as a strong sparsity constraint on the convolutional kernel as well as a drastic reduction of the number of learnable parameters. We construct a temporal

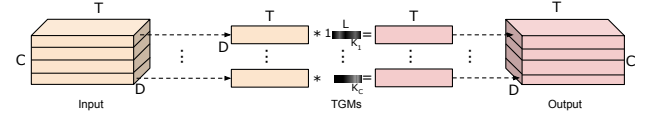


Figure 3. Illustration of a TGM layer with grouped convolution. This layer learns a set of C Gaussian mixtures that are convolved with the input channels.

Gaussian mixture convolutional kernel as:

$$\hat{K}_{m,l} = \frac{1}{Z} \exp - \frac{(l - \mu_m)^2}{2\sigma_m^2} \quad (2)$$

where Z is a normalization constant such that $\sum_l^L \hat{K}_{m,l} = 1$, resulting in \hat{K} being an $M \times L$ matrix.

Instead of making the model learn a separate set of Gaussian distributions per activity class, we take the approach of maintaining multiple Gaussian distributions shared across classes and obtain a Gaussian ‘mixture’ filter by learning soft-attention weights. We learn a set of soft-attention weights per output channel i , $\omega \in \mathcal{R}^{C_{out} \times M}$. We create the soft-attention weights by applying the softmax function over the M Gaussians, enforcing each input channel weights sum to 1.

$$a_{i,m} = \frac{\exp \omega_{i,m}}{\sum_j \exp \omega_{i,j}} \quad (3)$$

Based on temporal Gaussian distributions \hat{K}_i and attention weights $a_{i,m}$, the temporal convolution filters our TGM layer is computed as:

$$K_i = \sum_m a_{i,m} \hat{K}_i. \quad (4)$$

This provides us convolutional filters having the form of a mixture of temporal Gaussians, controlled based on $2 \cdot M + C_{in} \cdot C_{out} \cdot M$ parameters (instead of learning $D^2 \cdot L$ parameters without any constraint, as in standard temporal convolution where $C \ll D$). An overview of this process is shown in Fig. 1.

3.1.1. SINGLE TGM LAYER - DIRECT PER-CLASS ACTIVITY MODELING

The representation we obtain by applying our base CNNs to each frame (or local segment) has the dimensionality of D , and stacking them along time axis provides us the representation with $1 \times D \times T$ -dim. That is, in the case of using only one TGM layer to capture activity representations, our C_{in} is fixed to 1 and C_{out} is fixed to be the number of activity classes. This is the simplest case of our model, attaching one TGM layer on top of the $1 \times D \times T$ representation.

Our convolutional kernel, K , has a learned Gaussian mixture for each activity class. Let the video features v be a $D \times T$ matrix. Each K_i is a 2-D convolutional filter with

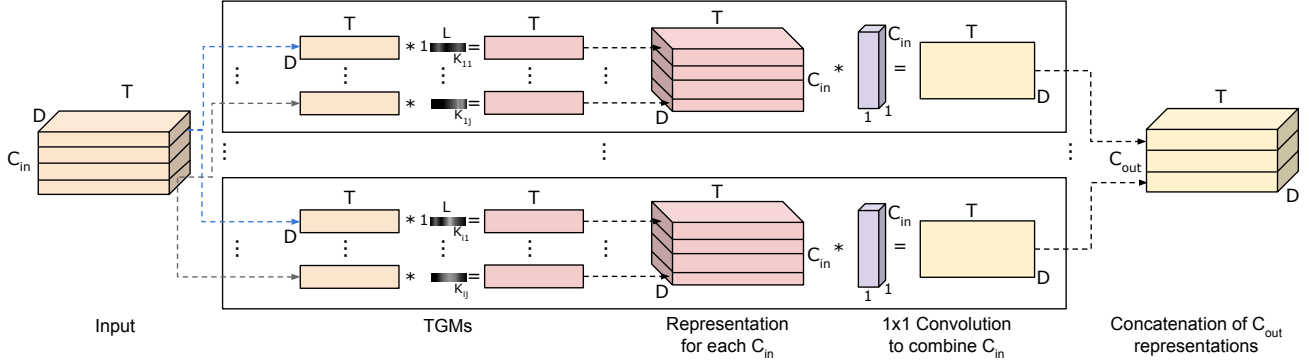


Figure 4. Illustration of a TGM layer with channel combination. The kernels are applied to each input channel, C_{in} , and a 1×1 convolution is applied to combine the C_{in} input channels for each output channel, C_{out} .

a size of $1 \times L$, and convolving this with v provides us a representation S with C_{out} number of $D \times T$ responses since C_{in} is 1 in this case. This per-class representation can then be used as input to a fully-connected layer for activity classification. For $i \in \{1, 2, \dots, C_{out}\}$:

$$s_i = v * K_i, \quad S = [s_1, s_2, \dots, s_{C_{out}}] \quad (5)$$

Fig. 2 visually illustrates how each TGM filter is convolved with the input (Fig. 2d), compared to the standard 1-D convolution (Fig. 2a) or other forms of the temporal layers (Fig. 2b-c).

3.1.2. MULTIPLE TGM LAYERS - GROUPED CONVOLUTION

We generalize the above formulation to allow the TGM layers to be sequentially applied. The idea is to enable our model to capture more complex, nonlinear temporal structure by having multiple levels of temporal layers. In this case, the input for each layer is $C_{in} \times D \times T$ dimensional (instead of $1 \times D \times T$), where the input channels are the number of output channels from the previous layer. Our kernels at each layer, K_i , are parameterized and learned as before.

By using grouped convolution with the number of groups set to C_{in} , we can efficiently separate the input into per-channel values and convolve each of them with the designated K_i kernel, as shown in Fig. 3. That is, we learn a filter K_i per channel by setting $C_{in} = C_{out}$. For $i \in [1, C_{out}]$,

$$s_i = f_i * K_i, \quad S = [s_1, s_2, \dots, s_{C_{out}}] \quad (6)$$

Here, f is a $C_{in} \times D \times T$ tensor, where D is the dimensionality of the feature and T is the number of frames. The result of the per-channel convolution, s_i , is a $D \times T$ representation. We concatenate these representations along the channel axis, resulting in S , a $C_{out} \times D \times T$ representation. As this convolution results in the same output shape, we can

stack these layers. Each layer is able to capture increasing temporal resolution, allowing the model to capture levels of abstractions.

3.1.3. MULTIPLE TGM LAYERS - CHANNEL COMBINATION

In the above subsection, we introduced an approach of stacking multiple TGM layers to model a hierarchical composition of temporal representations. However, in the grouped convolution case, each output channel of the layer is solely dependent on its corresponding input channel. That is, each kernel only considers information from a single output channel of the previous layer.

Therefore, we further generalize our TGM layer so that the layer combines representations from multiple input channels for each output channel while using the learned temporal kernels. We learn a set of convolutional kernels $K \in \mathcal{R}^{C_{out} \times C_{in} \times L}$ (i.e., we learn $C_{out} \cdot C_{in}$ Gaussian mixtures). Given f which is the $C_{in} \times D \times T$ representation, for each output channel $i \in [1, C_{out}]$ and each input channel $j \in [1, C_{in}]$ pair, we convolve the associated filters with the input.

$$G_{i,j} = (f_j * K_{i,j}) \quad (7)$$

where each $G_{i,j}$ is a $D \times T$ -dim representation.

We then learn a 1×1 convolution followed by a ReLU activation function for each $i \in [1, C_{out}]$, which we call w_i , that maps from C_{in} channels to 1 channel. The 1×1 convolution learns to combine the channels from the previous layer. By design, the TGM kernel is positive and sums to 1. Adding the unconstrained 1×1 convolution adds non-linearity (using the ReLU activation function) to our layer while only adding $C_{out} \cdot C_{in}$ parameters. The layer is computed as:

$$s_i = G_i * w_i = (f_j * K_{i,j}) * w_i, \quad S = [s_1, s_2, \dots, s_{C_{out}}] \quad (8)$$

We then stack the s_i representations along the channel axis to produce S , the $C_{out} \times D \times T$ -dim representation. This

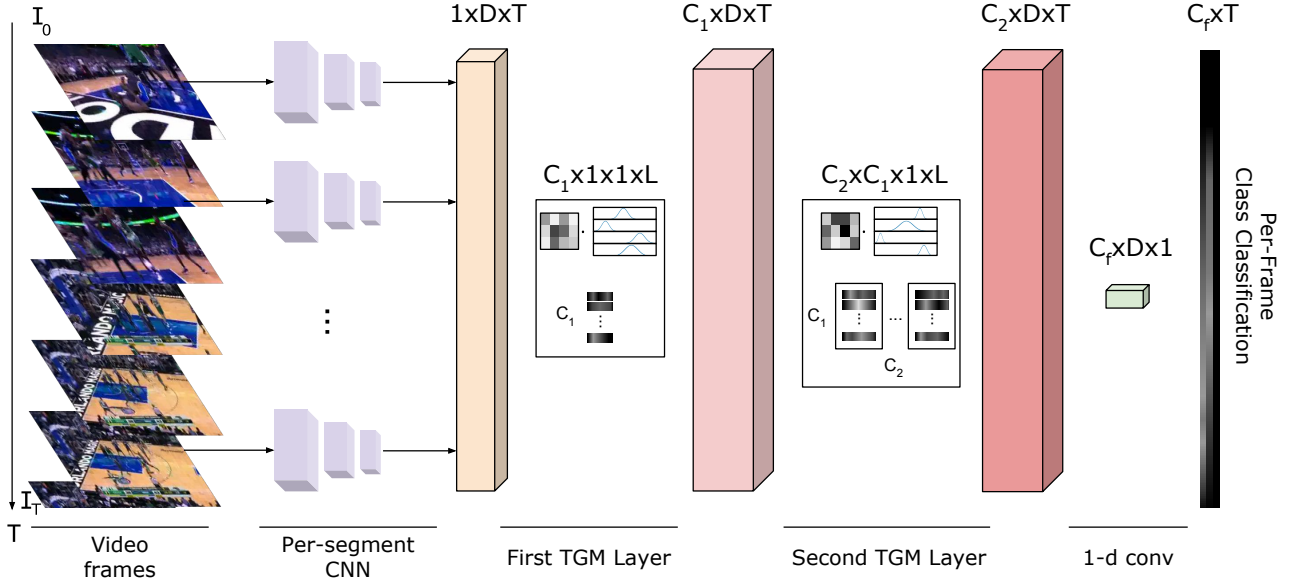


Figure 5. An overview of an example video CNN model with two TGM layers. Because of its fully convolutional design, it is able to handle videos with any length.

process is illustrated in Fig. 4. This method generalizes our approach to allow the layer to take input of $C_{in} \times D \times T$ and produce output of $C_{out} \times D \times T$. These layers can easily be stacked to learn a hierarchical representation.

3.2. Video CNN models with TGM layers

Our goal is to do activity detection which we define as making a per-frame (or per-segment) classification. Given a video, at each time step t , we want to make the model decide which activity the frame corresponds to (including no-activity). As a baseline, we train a fully-connected layer that classifies each per-frame D -dimensional vector, v_t . As multiple activities can occur at the same time, or no activities at all, we treat this as a multi-label classification task. We minimize binary cross entropy:

$$L(v) = \sum_{t,c} z_{t,c} \log(p(c|v_t)) + (1 - z_{t,c}) \log(1 - p(c|v_t)) \quad (9)$$

where $z_{t,c}$ is the ground truth label, 1 if activity c is occurring at time t and $p(c|v_t)$ is the output of our model for class c at time t . Fig. 5 shows an example CNN.

4. Experiments

4.1. Implementation and baselines

Implementation We used I3D (Carreira & Zisserman, 2017) and the two-stream version of InceptionV3 (Szegedy et al., 2016) pretrained on Imagenet and Kinetics as our base per-frame CNNs. Our default L setting used for the TGM layers as well as the other baselines was as follows:

when using I3D segment features (3 features per second from 24fps videos), the 1 layer models used $L = 15$ and the 3 layer models used $L = 5$. When using InceptionV3 frame feature (at 8 fps), the 1 layer models used $L = 30$ and the 3 layer models used $L = 10$. These layers were attached on top of the base CNN, as described in Subsection 3.2. Please check Appendix for implementation and training details and results on other datasets.

Baselines In order to confirm the advantages of our TGM layers, particularly against previous temporal models, we implemented many baselines. The first is (i) a standard per-frame classifier in which the prediction at each time-step only depends on a single feature vector with no contextual temporal information. We also used (ii) LSTMs on top of per-frame representations, which were popularly used to capture temporal information (Donahue et al., 2015). We train a bi-directional LSTM with 512 hidden units to make per-frame predictions. We also tried (iii) the fixed pyramid temporal max-pooling of level 3 (Ryoo et al., 2015).

We also compare our model against (iv) the model with standard temporal convolutional layers (i.e., 1-D convolution with a $D \times L$ kernel) on top of per-frame representations. This is similar to the temporal conv. used in (Tran et al., 2018). Temporal lengths (i.e., L) of the 1-D conv. filters and the pooling windows were set to be identical to the TGM filters. That is, they capture the same temporal duration as TGMs. In all our experiments, we follow the standard evaluation setting of computing per-frame mean average precision (mAP) and report those values.

Table 1. Comparison of various architectures on MultiTHUMOS using both I3D per-segment and InceptionV3 per-frame features. We found that TGM layers with 1x1 convolution channel combination performed the best. Results are in mAP. Note that we use the same filter length for “Temporal Conv” and “TGM” models, as described in Section 4.1.

	I3D			InceptionV3		
	Spatial	Temporal	Two-Stream	Spatial	Temporal	Two-Stream
Baseline	22.3	25.0	29.7	13.6	14.1	15.2
Temporal Conv	32.5	35.5	38.4	15.2	15.5	15.8
3 Temporal Conv	20.4	23.4	24.4	5.3	6.1	6.5
TGM layers with grouped convolution						
1 TGM	35.1	37.8	40.5	16.3	17.5	18.0
3 TGM	36.4	42.3	43.5	17.5	18.3	19.2
TGM layers with channel combination						
1 TGM (soft)	35.2	37.9	40.2	17.2	17.6	18.4
1 TGM (1x1)	36.1	38.2	40.8	17.2	17.7	18.4
3 TGM (soft)	36.2	40.1	42.3	17.5	19.1	21.2
3 TGM (1x1)	37.2	42.1	44.3	17.9	19.3	22.2
5 TGM (1x1)	37.4	42.4	44.8	18.2	20.3	23.4

We also compare to many different forms of the TGM layer: (v) 1-D convolution with a single Gaussian mixture per input (Fig. 2b); (vi) 1-D convolution with the kernel consisting of many Gaussian mixtures (Fig. 2c); (vii) our TC-grouping with unconstrained kernels (see Appendix for figure); (viii) with a learned mixture of random temporal filters and (ix) with a learned mixture of fixed Gaussians.

In addition, we also tried the approach of combining our TGM layers with the super-event representations, which capture global context (Piergiovanni & Ryoo, 2018b). We concatenated the learned super-event representation with our representations from TGM layers.

4.2. THUMOS / MultiTHUMOS

Dataset We conducted our experiments on both THUMOS (Jiang et al., 2014) and MultiTHUMOS (Yeung et al., 2015) datasets, while using the more challenging MultiTHUMOS as the main dataset. MultiTHUMOS is an extended version of the THUMOS dataset that densely annotates the continuous videos. The dataset consists of 65 different classes, compared to 20 in THUMOS, and contains on average 10.5 activities per video and 1.5 labels per frame and up to 25 activity instances in each video. This is in contrast to many other activity detection datasets such as ActivityNet (Heilbron et al., 2015), which only has on average ~ 1 activity per video. THUMOS and MultiTHUMOS consists of YouTube videos of various sport activities like basketball/volleyball games, weight lifting, and track/field.

We followed the standard evaluation setting for each dataset (e.g., measuring mAP based on per-frame annotations in MultiTHUMOS). There are 1010 validation videos and 1574 test videos. We used these continuous validation videos for the training of our models. We did not take ad-

Table 2. Additional number of parameters for models when added to the base architecture (e.g., I3D or Inception V3).

Model	# of parameters
LSTM	10.5M
1 Temporal Conv	10.5M
3 Temporal Conv	31.5M
1 TGM Layer	10K
3 TGM Layers	100K
5 TGM Layers	200K

vantage of the separate training set with segmented videos; even without them, we outperformed the state-of-the-arts.

Results We compared baselines as well as multiple different versions of our architectures, shown in Table 1. The model with our TGM layers consistently outperformed baseline I3D (or InceptionV3) while using the same per-segment representations. Learning 3 TGM layers further improved the performance. We also note that the use of a 1x1 convolution + ReLU non-linearity improves performance more than soft-attention (weighted averaging), confirming that this additional non-linearity is beneficial. On the other hand, we found that stacking multiple standard temporal convolutional layers does not improve performance, often performing worse than the baseline. While a single standard temporal conv. layer improves over the baseline, stacking multiple layers significantly increases the number of parameters to learn (Table 2) and we suspect this causes overfitting with the limited amount of samples in the dataset.

In Table 3, we explicitly compare the results of using an LSTM or temporal conv. with a similar number of parameters to our TGM. This was done by making their temporal conv. filters to share values across multiple channels. These

Table 3. Comparison of previous methods with comparable number of parameters and random forms of our TGM layer (using two-stream I3D on MultiTHUMOS).

Model	mAP
LSTM with 100k parameters	6.5
Temporal Conv. with 100k parameters	7.3
TGM with random temporal filters	34.5
TGM with fixed Gaussians	38.5
Full TGM	44.3

Table 4. Comparison of the different forms of temporal convolution on MultiTHUMOS using RGB I3D features. We set $L = 15$ and used 1 layer models for these experiments.

Standard 1-D Convolution (Fig. 2a)	32.5
1-D Conv with shared Gaussian mixture (Fig. 2b)	28.6
1-D Conv with Gaussian mixtures (Fig. 2c)	33.2
TC-grouping with unconstrained kernel	32.8
Our TGM Layer	36.1

models result in nearly random performance, as they were not designed to cope with a small number of parameters. We also show results with a mixture of random (fixed) temporal filters and with a mixture of fixed Gaussians. These results confirm that (i) modeling the temporal structure as a learned Gaussian mixture is beneficial and that (ii) further learning the Gaussian distribution parameters is important.

Also, in Table 4, we compare to the different forms of temporal convolution (as illustrated in Fig. 2). We find that each filter using one Gaussian mixture across all channels is not beneficial, while using a Gaussian mixture per-input channel (i.e., standard 1-D conv. with Gaussian mixture kernels) outperforms standard 1-D conv. Further, we find that using TC-grouping outperforms standard 1-D conv even with unconstrained kernels, although this itself is not as effective as 1-D conv. with Gaussian mixtures. Finally, we find that our designed TGM layer performs the best, confirming that both modeling temporal information as Gaussian mixtures and our designed TC-grouping are useful for activity detection.

Learning multiple TGM layers with channel combination outperforms the grouped convolution version of TGM and all the baselines. We also experimented with a version using soft-attention weights to combine the TGM layer channels, in addition to our method (Fig. 4) of using 1×1 convolution followed by a ReLU (to gain non-linearity). We found that the 1×1 convolution performed better. We tested various number of Gaussian mixtures (i.e., output channels) and found that using 80 for the first and second layer and using 65 (i.e., number of classes) for the final layer performs best.

Tables 5 and 6 compare our model using TGM layers with multiple previous state-of-the-art approaches and baselines. Our approach meaningfully outperforms all previous ap-

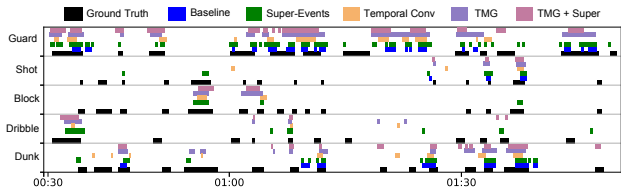


Figure 6. Illustration of the temporal regions classified as various basketball activities from a basketball game video in MultiTHUMOS. The TGM layers greatly improve performance.

Table 5. Performances of the state-of-the-art methods and our approach on MultiTHUMOS. Our approach meaningfully outperforms all previous results.

	mAP
Two-stream (Yeung et al., 2015)	27.6
Two-stream + LSTM (Yeung et al., 2015)	28.1
Multi-LSTM (Yeung et al., 2015)	29.6
Predictive-corrective (Dave et al., 2017)	29.7
SSN (Zhao et al., 2017)	30.3
I3D baseline	29.7
I3D + LSTM	29.9
I3D + temporal pyramid	31.2
I3D + super-events (Piergiovanni & Ryoo, 2018b)	36.4
I3D + our TGMs	44.3
I3D + super-events + our TGMs	46.4

proaches. Importantly, we are comparing our approach with different methods of capturing temporal information such as LSTMs and fixed temporal pyramid pooling while making them use the exactly same per-frame representations. We found that while all these methods capture some temporal information, the TGM layers provide the best performance. Further, combining the super-event representation (Piergiovanni & Ryoo, 2018b) with our TGM feature also benefited detection, confirming that the TGMs and super-events capture different aspects of the activity videos. In Fig. 6, we show an example of the various models predictions on a basketball video. We outperform the previous state-of-the-art performance (mAP) by 10% (36.4 vs. 46.4).

Table 6. Performances of our approach compared to the state-of-the-arts on the continuous version of THUMOS-14 dataset.

	mAP
R-C3D (Xu et al., 2017)	28.9
SSN (Zhao et al., 2017)	29.1
TAL-Net (Chao et al., 2018)	42.8
I3D Baseline	43.5
I3D + super-events (Piergiovanni & Ryoo, 2018b)	47.8
I3D + 1 TGMs	45.2
I3D + 3 TGMs	53.5
I3D + 3 TGMs + super-events	57.0

Table 7. Per-frame mAP on Charades, evaluated with the ‘Charades_v1_localize’ setting. I3D models are two-stream, using both RGB and optical flow inputs.

	mAP
Predictive-corrective (Dave et al., 2017)	8.9
Two-stream (Sigurdsson et al., 2016a)	8.94
Two-stream+LSTM (Sigurdsson et al., 2016a)	9.6
R-C3D (Xu et al., 2017)	12.7
Sigurdsson et al. (Sigurdsson et al., 2016a)	12.8
SSN (Zhao et al., 2017)	16.4
I3D baseline	17.2
I3D + 3 temporal conv. layers ($L = 5$)	17.5
I3D + 3 temporal conv. layers ($L = 30$)	12.5
I3D + LSTM	18.1
I3D + fixed temporal pyramid	18.2
I3D + super-events (Piergiovanni & Ryoo, 2018b)	19.4
I3D + 3 TGMs ($L = 5$)	20.6
I3D + 3 TGMs ($L = 30$)	21.5
I3D + 3 TGMs ($L = 5$) + super-events	21.8
I3D + 3 TGMs ($L = 30$) + super-events	22.3

4.3. Charades

Dataset Charades (Sigurdsson et al., 2016b) is a large scale dataset with 9848 videos across 157 activity classes. These videos were recorded in home environments of the participants based on provided scripts. Each video contains on an average of 6.8 activity instances, and there are often complex activities co-occurring. The activities were mainly performed at home. For example, some activity classes are ‘preparing a meal’, ‘eating’, ‘sitting’, ‘cleaning’, etc.

In our experiments, we follow the original Charades detection setting (i.e., Charades_v1_localize evaluation), which is the original setting used in many previous approaches (Sigurdsson et al., 2016a; Piergiovanni & Ryoo, 2018b).

Results We compare our results with the state-of-the-arts in Table 7. To our knowledge, our method is obtaining the best known performance in the original localization setting of the Charades dataset. Notably, it is performing better than I3D that obtained the best competition performance, while using the same feature. Our method also outperforms standard temporal convolution, LSTMs, and fixed pyramid pooling, as well as the use of latent super-events. When setting $L = 30$ and using 3 TGM layers, our model is able to capture around 800 frames (about ± 15 seconds from each frame) of temporal information, significantly more than previous works (e.g., I3D only captures ± 2 seconds). This confirms that a key advantage of the TGM layer is that the number of parameters is independent of the filter length.

In Table 8 (and Table 3 in Appendix), we compare different values of L . Due to the short duration of activities in MultiTHUMOS (3.3 seconds), we find that $L = 5$ with 3 layers



Figure 7. Illustration of several learned TGM kernels. On MultiTHUMOS, it learns to focus on shorter intervals to capture shorter events. On Charades, the Gaussians have a larger σ value, resulting in filters that attend to longer temporal durations.

Table 8. Effect of L on MultiTHUMOS using only RGB I3D features. Note that the 3 TGM layer models capture information in larger temporal intervals than the 1 TGM layer models for the same values of L . We also compare to using standard one-layer 1-D conv layer with different values of L .

	1 TGM	3 TGM	1-D Conv
I3D Baseline	22.3	-	-
$L = 3$	30.2	31.7	26.6
$L = 5$	32.5	37.2	28.3
$L = 10$	34.5	35.4	31.7
$L = 15$	36.1	34.1	32.5
$L = 30$	32.5	33.9	26.5
$L = 50$	32.1	33.7	15.4

performs the best. Larger values of L capture too much unneeded temporal information, but due to the Gaussian structure, it does not drastically harm performance. Figure 7 shows that even with longer kernels, the Gaussians learn to focus mostly on the center of the interval and capture short intervals. Thus, having too long intervals does not drastically harm performance, which is in contrast to the standard 1-D convolution. Note that for Charades, the temporal kernels learned to capture much longer temporal durations, as the average activity in charades is 12.8 seconds and larger values of L are better.

5. Conclusions

We newly introduced the Temporal Gaussian Mixture (TGM) layer and demonstrated its effectiveness for multi-activity detection in continuous videos. Our layer is fully differentiable and trainable using standard backpropagation, designed to learn temporal structure. We were able to confirm that our layer performs superior to state-of-the-art methods on activity detection datasets including MultiTHUMOS and Charades, obtaining the best known performance. We also tested our approach with two more public video datasets, MLB-YouTube (Piergiovanni & Ryoo, 2018a) and AVA (Gu et al., 2017), and confirmed its advantage over the previous works in Appendix.

Acknowledgements

This work was supported in part by the National Science Foundation (IIS-1812943 and CNS-1814985), and by the ICT R&D program of MSIP/IITP, Republic of Korea (17ZF1200, Development of XDMedia Solution for Invigoration of Realistic Media Industry).

References

- Aggarwal, J. K. and Ryoo, M. S. Human activity analysis: A review. *ACM Computing Surveys*, 43:16:1–16:43, April 2011.
- Carreira, J. and Zisserman, A. Quo vadis, action recognition? a new model and the kinetics dataset. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- Chao, Y.-W., Vijayanarasimhan, S., Seybold, B., Ross, D. A., Deng, J., and Sukthankar, R. Rethinking the faster r-cnn architecture for temporal action localization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 1130–1139, 2018.
- Cohen, T. S., Geiger, M., Köhler, J., and Welling, M. Spherical cnns. *arXiv preprint arXiv:1801.10130*, 2018.
- Dave, A., Russakovsky, O., and Ramanan, D. Predictive-corrective networks for action detection. *arXiv preprint arXiv:1704.03615*, 2017.
- Donahue, J., Anne Hendricks, L., Guadarrama, S., Rohrbach, M., Venugopalan, S., Saenko, K., and Darrell, T. Long-term recurrent convolutional networks for visual recognition and description. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 2625–2634, 2015.
- Escorcia, V., Heilbron, F. C., Niebles, J. C., and Ghanem, B. Daps: Deep action proposals for action understanding. In *Proceedings of European Conference on Computer Vision (ECCV)*, pp. 768–784. Springer, 2016.
- Feichtenhofer, C., Pinz, A., and Zisserman, A. Convolutional two-stream network fusion for video action recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 1933–1941, 2016.
- Gu, C., Sun, C., Vijayanarasimhan, S., Pantofaru, C., Ross, D. A., Toderici, G., Li, Y., Ricco, S., Sukthankar, R., Schmid, C., and Malik, J. AVA: A video dataset of spatio-temporally localized atomic visual actions. *arXiv preprint arXiv:1705.08421*, 2017.
- Hara, K., Kataoka, H., and Satoh, Y. Learning spatio-temporal features with 3d residual networks for action recognition. In *Proceedings of the ICCV Workshop on Action, Gesture, and Emotion Recognition*, volume 2, pp. 4, 2017.
- Heilbron, F. C., Escorcia, V., Ghanem, B., and Niebles, J. C. Activitynet: A large-scale video benchmark for human activity understanding. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 961–970, 2015.
- Jiang, Y.-G., Liu, J., Roshan Zamir, A., Toderici, G., Laptev, I., Shah, M., and Sukthankar, R. THUMOS challenge: Action recognition with a large number of classes. <http://crcv.ucf.edu/THUMOS14/>, 2014.
- Karpathy, A., Toderici, G., Shetty, S., Leung, T., Sukthankar, R., and Fei-Fei, L. Large-scale video classification with convolutional neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 1725–1732, 2014.
- Kay, W., Carreira, J., Simonyan, K., Zhang, B., Hillier, C., Vijayanarasimhan, S., Viola, F., Green, T., Back, T., Natsev, P., et al. The kinetics human action video dataset. *arXiv preprint arXiv:1705.06950*, 2017.
- Ng, J. Y.-H., Hausknecht, M., Vijayanarasimhan, S., Vinyals, O., Monga, R., and Toderici, G. Beyond short snippets: Deep networks for video classification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 4694–4702. IEEE, 2015.
- Piergiovanni, A. and Ryoo, M. S. Fine-grained activity recognition in baseball videos. In *CVPR Workshop on Computer Vision in Sports*, 2018a.
- Piergiovanni, A. and Ryoo, M. S. Learning latent super-events to detect multiple activities in videos. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018b.
- Piergiovanni, A., Fan, C., and Ryoo, M. S. Learning latent sub-events in activity videos using temporal attention filters. In *Proceedings of the American Association for Artificial Intelligence (AAAI)*, 2017.
- Ryoo, M. S., Rothrock, B., and Matthies, L. Pooled motion features for first-person videos. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 896–904, 2015.
- Shou, Z., Wang, D., and Chang, S.-F. Temporal action localization in untrimmed videos via multi-stage cnns. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 1049–1058, 2016.

- Shou, Z., Chan, J., Zareian, A., Miyazawa, K., and Chang, S.-F. Cdc: Convolutional-de-convolutional networks for precise temporal action localization in untrimmed videos. *arXiv preprint arXiv:1703.01515*, 2017.
- Sigurdsson, G. A., Divvala, S., Farhadi, A., and Gupta, A. Asynchronous temporal fields for action recognition. *arXiv preprint arXiv:1612.06371*, 2016a.
- Sigurdsson, G. A., Varol, G., Wang, X., Farhadi, A., Laptev, I., and Gupta, A. Hollywood in homes: Crowdsourcing data collection for activity understanding. In *Proceedings of European Conference on Computer Vision (ECCV)*, 2016b.
- Simonyan, K. and Zisserman, A. Two-stream convolutional networks for action recognition in videos. In *Advances in Neural Information Processing Systems (NIPS)*, pp. 568–576, 2014.
- Szegedy, C., Vanhoucke, V., Ioffe, S., Shlens, J., and Wojna, Z. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 2818–2826, 2016.
- Tran, D., Bourdev, L. D., Fergus, R., Torresani, L., and Paluri, M. C3d: generic features for video analysis. *CoRR*, *abs/1412.0767*, 2(7):8, 2014.
- Tran, D., Ray, J., Shou, Z., Chang, S.-F., and Paluri, M. Convnet architecture search for spatiotemporal feature learning. *arXiv preprint arXiv:1708.05038*, 2017.
- Tran, D., Wang, H., Torresani, L., Ray, J., LeCun, Y., and Paluri, M. A closer look at spatiotemporal convolutions for action recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 6450–6459, 2018.
- Variani, E., McDermott, E., and Heigold, G. A gaussian mixture model layer jointly optimized with discriminative features within a deep neural network architecture. In *Acoustics, Speech and Signal Processing (ICASSP), 2015 IEEE International Conference on*, pp. 4270–4274. IEEE, 2015.
- Varol, G., Laptev, I., and Schmid, C. Long-term Temporal Convolutions for Action Recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2017.
- Xu, H., Das, A., and Saenko, K. R-c3d: Region convolutional 3d network for temporal activity detection. *arXiv preprint arXiv:1703.07814*, 2017.
- Xu, Y., Fan, T., Xu, M., Zeng, L., and Qiao, Y. Spidercnn: Deep learning on point sets with parameterized convolutional filters. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pp. 87–102, 2018.
- Yeung, S., Russakovsky, O., Jin, N., Andriluka, M., Mori, G., and Fei-Fei, L. Every moment counts: Dense detailed labeling of actions in complex videos. *International Journal of Computer Vision (IJCV)*, pp. 1–15, 2015.
- Yeung, S., Russakovsky, O., Mori, G., and Fei-Fei, L. End-to-end learning of action detection from frame glimpses in videos. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 2678–2687, 2016.
- Zhao, Y., Xiong, Y., Wang, L., Wu, Z., Tang, X., and Lin, D. Temporal action detection with structured segment networks. *arXiv preprint arXiv:1704.06228*, 2017.