

FIFA: Fast Inference Approximation for Action Segmentation

Yaser Souri¹, Yazan Abu Farha¹, Fabien Despinoy², Gianpiero Francesca², and Juergen Gall¹

¹ University of Bonn

² Toyota Motor Europe

{souri, abufarha, gall}@iai.uni-bonn.de

{fabien.despinoy, gianpiero.francesca}@toyota-motor.com

Abstract. We introduce FIFA, a fast approximate inference method for action segmentation and alignment. Unlike previous approaches, FIFA does not rely on expensive dynamic programming for inference. Instead, it uses an approximate differentiable energy function that can be minimized using gradient-descent. FIFA is a general approach that can replace exact inference, improving its speed by more than 5 times while maintaining its performance. FIFA is an anytime inference algorithm that provides a better speed vs. accuracy trade-off compared to exact inference. We apply FIFA on top of state-of-the-art approaches for weakly supervised action segmentation and alignment as well as fully supervised action segmentation. FIFA achieves state-of-the-art results for most metrics on two action segmentation datasets.

Keywords: Video Understanding · Action Segmentation · Approximate Inference.

1 Introduction

Action segmentation is the task of predicting the action label for each frame in the input video. Action segmentation is usually studied in the context of activities performed by a single person, where temporal smoothness of actions are assumed. Fully supervised approaches for action segmentation [23,1,28,36] already achieve good performance on this task. Most approaches for fully supervised action segmentation make frame-wise predictions [23,1,28] while trying to model the temporal relationship between the action labels. These approaches usually suffer from over-segmentation. Recent works [36,14] try to overcome the over-segmentation problem by finding the action boundaries and temporally smoothing the predictions inside each action segment. But these post-processing approaches still can not guarantee temporal smoothness.

Action segmentation inference is the problem of making segment-wise smooth predictions from frame-wise probabilities given a known grammar of the actions and their average lengths [32]. The typical inference in action segmentation involves solving an expensive Viterbi-like dynamic programming problem that

finds the best action sequence and its corresponding lengths. In the literature, weakly supervised action segmentation approaches [19,31,32,27,34] usually use inference at test time. Despite being very useful for action segmentation, the inference problem remains the main computational bottleneck in the action segmentation pipeline [34].

In this paper, we propose FIFA, a fast anytime approximate inference procedure that achieves comparable performance with respect to the dynamic programming based Viterbi decoding inference at a fraction of the computational time. Instead of relying on dynamic programming, we formulate the energy function as an approximate differentiable function of segment lengths parameters and use gradient-descent-based methods to search for a configuration that minimizes the approximate energy function. Given a transcript of actions and the corresponding initial lengths configuration, we define the energy function as a sum over segment level energies. The segment level energy consists of two terms: a length energy term that penalizes the deviations from a global length model and an observation energy term that measures the compatibility between the current configuration and the predicted frame-wise probabilities. A naive approach to model the observation energy would be to sum up the negative log probabilities of the action labels that are defined based on the length configuration. Nevertheless, such an approach is not differentiable with respect to the segment lengths. In order to optimize the energy using gradient descent-based methods, the observation energy has to be differentiable with respect to the segment lengths. To this end, we construct a plateau-shaped mask for each segment which temporally locates the segment within the video. This mask is parameterized by the segment lengths, the position in the video, and a sharpness parameter. The observation energy is then defined as a product of a segment mask and the predicted frame-wise negative log probabilities, followed by a sum pooling operation. Finally, a gradient descent-based method is used to find a configuration for the segment lengths that minimizes the total energy.

FIFA is a general inference approach and can be applied at test time on top of different action segmentation approaches for fast inference. We evaluate our approach on top of the state-of-the-art methods for weakly supervised temporal action segmentation, weakly supervised action alignment, and fully supervised action segmentation. Results on the Breakfast [17] and Hollywood extended [4] datasets show that FIFA achieves state-of-the-art results on most metrics. Compared to the exact inference using the Viterbi decoding, FIFA is at least 5 times faster. Furthermore, FIFA is an anytime algorithm which can be stopped after each step of the gradient-based optimization, therefore it provides a better speed vs. accuracy trade-off compared to exact inference.

2 Related Work

In this section we highlight relevant works addressing fully and weakly supervised action segmentation that have been recently proposed.

Fully Supervised Action Segmentation. In fully supervised action segmentation, frame-level labels are used for training. Initial attempts for action segmentation applied action classifiers on a sliding window over the video frames [33,16]. However, these approaches did not capture the dependencies between the action segments. With the objective of capturing the context over long video sequences, context free grammars [35,30] or hidden Markov models (HMMs) [24,18,21] are typically combined with frame-wise classifiers. Recently, temporal convolutional networks showed good performance for the temporal action segmentation task using encoder-decoder architectures [23,26] or even multi-stage architectures [1,28]. Many approaches further improve the multi-stage architectures by applying post-processing based on boundary-aware pooling operation [36,14] or graph-based reasoning [13]. Without any inference most of the fully-supervised approaches therefore suffer from oversegmentation at test time.

Weakly Supervised Action Segmentation. To reduce the annotation cost, many approaches that rely on a weaker form of supervision have been proposed. Earlier approaches apply discriminative clustering to align video frames to movie scripts [8]. Bojanowski *et al.* [5] proposed to use as supervision the transcripts in the form of ordered lists of actions. Indeed, many approaches rely on this form of supervision to train a segmentation model using connectionist temporal classification [12], dynamic time warping [6] or energy-based learning [27]. In [7], an iterative training procedure is used to refine the transcript. A soft labeling mechanism is further applied at the boundaries between action segments. Kuehne *et al.* [20] applied a speech recognition system based on a HMM and Gaussian mixture model (GMM) to align video frames to transcripts. The approach generates pseudo ground truth labels for the training videos and iteratively refines them. A similar idea has been recently used in [31,21]. Richard *et al.* [32] combined the frame-wise loss function with the Viterbi algorithm to generate the target labels. At inference time, these approaches iterate over the training transcripts and select the one that matches the test video best. By contrast, Souri *et al.* [34] predict the transcript besides the frame-wise scores at inference time. State-of-the-art weakly supervised action segmentation approaches require time consuming dynamic programming based inference at test time.

Energy-Based Inference. In energy-based inference methods, gradient descent is used at inference time as described in [25]. The goal is to minimize an energy function that measures the compatibility between the input variables and the predicted variables. This idea has been exploited for many structured prediction tasks such as image generation [9,15], machine translation [11] and structured prediction energy networks [3]. Belanger and McCallum [2] relaxed the discrete output space for multi-label classification tasks to a continuous space and used gradient descent to approximate the solution. Gradient-based methods have also been used for other applications such as generating adversarial examples [10] and learning text embeddings [22].

3 Background

The following sections introduce all the concepts and notations required to understand the proposed FIFA methodology.

3.1 Action Segmentation

In action segmentation, we want to temporally localize all the action segments occurring in a video. In this paper, we consider the case where the actions are from a predefined set of M classes (a background class is used to cover uninteresting parts of a video). The input video of length T is usually represented as a set of d dimensional features vectors $x_{1:T} = (x_1, \dots, x_T)$. These features are extracted offline and are assumed to be the input to the action segmentation model. The output of action segmentation can be represented in two ways:

- Frame-wise representation $y_{1:T} = (y_1, \dots, y_T)$ where y_t represents the action label at time t .
- Segment-wise representation $s_{1:N} = (s_1, \dots, s_N)$ where segment s_n is represented by both the action label of the segment c_n and its corresponding length ℓ_n , i.e., $s_n = (c_n, \ell_n)$. The ordered list of actions $c_{1:N}$ is usually referred to as the *transcript*.

These two representations are equal and redundant, i.e., it is possible to compute one from the other. In order to transfer from the segment-wise to the frame-wise representation, we introduce a mapping $\alpha(t; c_{1:N}, \ell_{1:N})$ which outputs the action label at frame t given the segment-wise labeling.

The target labels to train a segmentation model, depend on the level of supervision. In fully supervised action segmentation [1,28,36], the target label for each frame is provided. However, in weakly supervised approaches [32,27,34] only the ordered list of action labels are provided during training while their lengths are unknown.

Recent fully supervised approaches for action segmentation like MSTCN [1] and its variants directly predict the frame-wise representation $y_{1:T}$ by choosing the action label with the highest probability for each frame independently. This results in predictions that are sometimes oversegmented.

Conversely, recent weakly supervised action segmentation approaches like NNV [32] and follow-up work include an inference stage during testing where they explicitly predict the segment-wise representation. This inference stage involves a dynamic programming algorithm for solving an optimization problem which is a computational bottleneck for these approaches.

3.2 Inference in Action Segmentation

During testing, the inference stage involves an optimization problem to find the most likely segmentation for the input video, i.e.,

$$c_{1:N}, \ell_{1:N} = \underset{\hat{c}_{1:N}, \hat{\ell}_{1:N}}{\operatorname{argmax}} \left\{ p(\hat{c}_{1:N}, \hat{\ell}_{1:N} | x_{1:T}) \right\}. \quad (1)$$

Given the transcript $c_{1:N}$, the inference stage boils down to finding the segment lengths $\ell_{1:N}$ by aligning the transcript to the input video, i.e.,

$$\ell_{1:N} = \operatorname{argmax}_{\hat{\ell}_{1:N}} \left\{ p(\hat{\ell}_{1:N} | x_{1:T}, c_{1:N}) \right\}. \quad (2)$$

In approaches like NNV [32] and CDFL [27], the transcript is found by iterating over the transcripts seen during training and selecting the transcript that achieves the most likely alignment by optimizing (2). In MuCon [34], the transcript is predicted by a sequence to sequence network.

The probability defined in (2) is broken down by making independences assumption between frames

$$p(\hat{\ell}_{1:N} | x_{1:T}, c_{1:N}) = \prod_{t=1}^T p(\alpha(t; c_{1:N}, \hat{\ell}_{1:N}) | x_t) \cdot \prod_{n=1}^N p(\hat{\ell}_n | c_n) \quad (3)$$

where $p(\alpha(t) | x_t)$ is referred to as the observation model and $p(\ell_n | c_n)$ as the length model. Here, $\alpha(t)$ is the mapping from time t to the action label given the segment-wise labeling. The observation model estimates the frame-wise action probabilities and is implemented using a neural network. The length model is used to constrain the inference defined in (2) with the assumption that the lengths of segments for the same action follow a particular probability distribution. The segment length is usually modelled by a Poisson distribution with a class dependent mean parameter λ_{c_n} , i.e.,

$$p(\ell_n | c_n) = \frac{\lambda_{c_n}^{\ell_n} \exp(-\lambda_{c_n})}{\ell_n!}. \quad (4)$$

This optimization is solved using an expensive dynamic programming based Viterbi decoding [32]. For details on how to solve this optimization problem using Viterbi decoding please refer to the supplementary material.

4 FIFA: Fast Inference Approximation

Our goal is to introduce a fast inference algorithm for action segmentation. We want the fast inference to be applicable in both weakly supervised and fully supervised action segmentation. We also want the fast inference to be flexible enough to work with different action segmentation methods. To this end, we introduce FIFA, a novel approach for fast inference for action segmentation.

In the following, for brevity, we write the mapping $\alpha(t; c_{1:N}, \ell_{1:N})$ simply as $\alpha(t)$. Maximizing probability (2) can be rewritten as minimizing the negative logarithm of that probability

$$\operatorname{argmax}_{\hat{\ell}_{1:N}} \left\{ p(\hat{\ell}_{1:N} | x_{1:T}, c_{1:N}) \right\} = \operatorname{argmin} \left\{ -\log(p(\hat{\ell}_{1:N} | x_{1:T}, c_{1:N})) \right\} \quad (5)$$

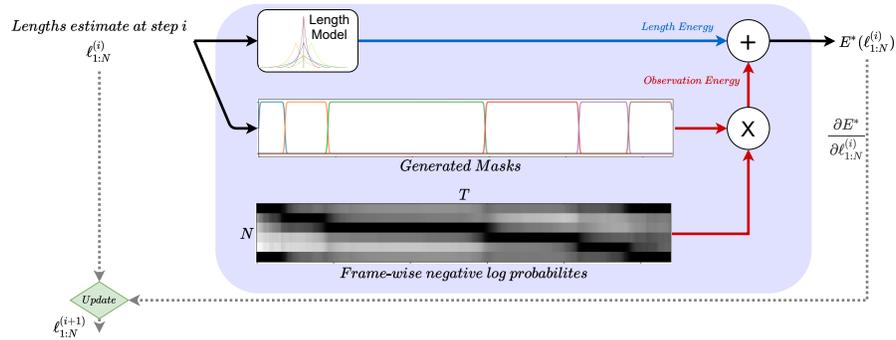


Fig. 1. Overview of the FIFA optimization process. At each step in the optimization, a set of masks are generated using the current length estimates. Using the generated masks and the frame-wise negative log probabilities, the observation energy is calculated in an approximate but differentiable manner. The length energy is calculated from the current length estimate and added to the observation energy to calculate the total energy value. Taking the gradient of the total energy with respect to the length estimates we can update it using a gradient step.

which we refer to as the energy $E(\ell_{1:N})$. Using (3) the energy is rewritten as

$$\begin{aligned}
 E(\ell_{1:N}) &= -\log \left(p(\ell_{1:N} | x_{1:T}, c_{1:N}) \right) \\
 &= -\log \left(\prod_{t=1}^T p(\alpha(t) | x_t) \cdot \prod_{n=1}^N p(\ell_n | c_n) \right) \\
 &= \underbrace{\sum_{t=1}^T -\log p(\alpha(t) | x_t)}_{E_o} + \underbrace{\sum_{n=1}^N -\log p(\ell_n | c_n)}_{E_\ell}.
 \end{aligned} \tag{6}$$

The first term in (6), E_o is referred to as the observation energy. This term calculates the cost of assigning the labels for each frame and is calculated from the frame-wise probability estimates. The second term E_ℓ is referred to as the length energy. This term is the cost of each segment having a length given that we assume an average length for actions of a specific class.

We propose to optimize the energy defined in (6) using gradient based optimization in order to avoid the need for time-consuming dynamic programming. We start with an initial estimate of the lengths (obtained from the length model of each approach or calculated from training data when available) and update our estimate to minimize the energy function.

As the energy function $E(\ell_{1:N})$ is not differentiable with respect to the lengths, we have to calculate a relaxed and approximate energy function $E^*(\ell_{1:N})$ that is differentiable.

4.1 Approximate Differentiable Energy E^*

The energy function E as defined in (6) is not differentiable in two parts. First the observation energy term E_o is not differentiable because of the $\alpha(t)$ function. Second, the length energy term E_ℓ is not differentiable because it expects natural numbers as input and cannot be computed on real values which are dealt with in gradient-based optimization. Below we describe how we approximate and make each of the terms differentiable.

Approximate Differentiable Observation Energy Consider a $N \times T$ matrix P containing negative log probabilities, i.e.,

$$P[n, t] = -\log p(c_n | x_t). \tag{7}$$

Furthermore, we define a mask matrix M with the same size $N \times T$ where

$$M[n, t] = \begin{cases} 0 & \text{if } \alpha(t) \neq c_n \\ 1 & \text{if } \alpha(t) = c_n \end{cases}. \tag{8}$$

Using the mask matrix we can rewrite the observation energy term as

$$E_o = \sum_{t=1}^T \sum_{n=1}^N M[n, t] \cdot P[n, t]. \tag{9}$$

In order to make the observation energy term differentiable with respect to the length, we propose to construct an approximate differentiable mask matrix M^* . We use the following smooth and parametric plateau function

$$f(t | \lambda^c, \lambda^w, \lambda^s) = \frac{1}{(e^{\lambda^s(t - \lambda^c - \lambda^w)} + 1)(e^{\lambda^s(-t + \lambda^c - \lambda^w)} + 1)} \tag{10}$$

from [29]. This plateau function has three parameters and it is differentiable with respect to them: λ^c controls the center of the plateau, λ^w is the width and λ^s is the sharpness of the plateau function.

While the sharpness of the plateau functions λ^s is fixed as a hyper-parameter of our approach, the center λ^c and the width λ^w are computed from the lengths $\ell_{1:N}$. First we calculate the starting position of each plateau function b_n as

$$b_1 = 0, b_n = \sum_{n'=1}^{n-1} \ell_{n'}. \tag{11}$$

We can then define both the center and the width parameters of each plateau function as

$$\begin{aligned} \lambda_n^c &= b_n + \ell_n/2, \\ \lambda_n^w &= \ell_n/2 \end{aligned} \tag{12}$$

and define each row of the approximate mask as

$$M^*[n, t] = f(t|\lambda_n^c, \lambda_n^w, \lambda^s). \quad (13)$$

Now we can calculate a differentiable approximate observation energy similar to (9) as

$$E_o^* = \sum_{t=1}^T \sum_{n=1}^N M^*[n, t] \cdot P[n, t]. \quad (14)$$

Approximate Differentiable Length Energy For the gradient-based optimization, we must relax the length values to be positive real values instead of natural numbers. As the Poisson distribution (4) is only defined on natural numbers, we propose to use a substitute distribution defined on real numbers. As a replacement, we experiment with a Laplace distribution and a Gaussian distribution. In both cases, the scale or the width parameter of the distribution is assumed to be fixed.

We can rewrite the length energy E_ℓ as the approximate length energy

$$E_\ell^*(\ell_{1:N}) = \sum_{n=1}^N -\log p(\ell_n | \lambda_{c_n}^\ell), \quad (15)$$

where $\lambda_{c_n}^\ell$ is the expected value for the length of a segment from the action c_n . In case of the Laplace distribution this length energy will be

$$E_\ell^*(\ell_{1:N}) = \frac{1}{Z} \sum_{n=1}^N |\ell_n - \lambda_{c_n}^\ell|, \quad (16)$$

where Z is the constant normalization factor. This means that the length energy will penalize any deviation from the expected average length linearly. Similarly, for the Gaussian distribution, the length energy will be

$$E_\ell^*(\ell_{1:N}) = \frac{1}{Z} \sum_{n=1}^N |\ell_n - \lambda_{c_n}^\ell|^2, \quad (17)$$

which means that the Gaussian length energy will penalize any deviation from the expected average length quadratically.

With the objective to maintain a positive value for the length during the optimization process, we estimate the length in log space and convert it to absolute space only in order to compute both the approximate mask matrix M^* and the approximate length energy E_ℓ^* .

Approximate Energy Optimization The total approximate energy function is defined as a weighted sum of both the approximate observation and the approximate length energy functions

$$E^*(\ell_{1:N}) = E_o^*(\ell_{1:N}, Y) + \beta E_\ell^*(\ell_{1:N}) \quad (18)$$

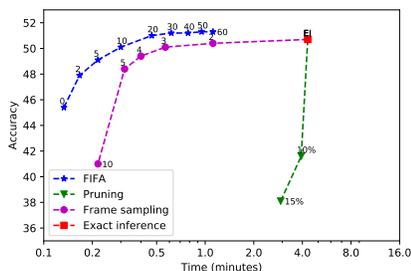


Fig. 2. Speed vs. accuracy trade-off of different inference approaches applied to the MuCon method. Using FIFA we can achieve a better speed vs. accuracy trade-off compared to frame sampling or hypothesis pruning in exact inference.

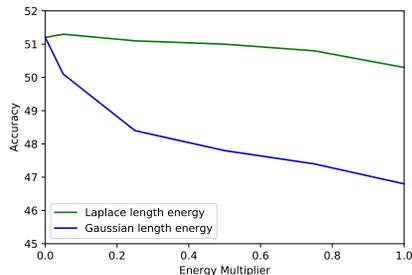


Fig. 3. Effect of the length energy multiplier for Laplace and Gaussian length energy. Accuracy is calculated on the Breakfast dataset using FIFA applied to the MuCon approach trained in the weakly supervised action segmentation setting.

where β is the multiplier for the length energy.

Given an initial length estimate $\ell_{1:N}^0$, we iteratively update this estimate to minimize the total energy. Figure 1 illustrates the optimization step for our approach. During each optimization step, we first calculate the energy E^* and then calculate the gradients of the energy with respect to the length values. Using the calculated gradients, we update the length estimate using a gradient descent update rule such as SGD or Adam. After a certain number of gradient steps (50 steps in our experiments) we will finally predict the segment length.

If the transcript for a test video is provided then it is used ,e.g., using the MuCon [34] approach or in a weakly supervised action alignment setting. However, if the latter is not known ,e.g., in a fully supervised approach or CDFL [27] for weakly supervised action segmentation, we perform the optimization for each of the transcripts seen during training and select the most likely one based on the final energy value at the end of the optimization.

The initial length estimates are calculated from the length model of each approach in case of a weakly supervised setting whereas in a fully supervised setting the average length of each action class is calculated from the training data and used as the initial length estimates. The initial length estimates are also used as the expected length parameters for the length energy calculations.

The hyper-parameters like the choice of the optimizer, number of steps, learning rate, and the mask sharpness, remain as the hyper-parameters of our approach.

5 Experiments

5.1 Evaluation Protocols and Datasets

We evaluate FIFA on 3 different tasks: weakly supervised action segmentation, fully supervised action segmentation, and weakly supervised action alignment.

Results for action alignment are included in the supplementary material. We obtain the source code for the state-of-the-art approaches on each of these tasks and train a model using the standard training configuration of each model. Then we apply FIFA as a replacement for an existing inference stage or as an additional inference stage.

We evaluate our model using the Breakfast [17] and Hollywood extended [4] datasets on the 3 different tasks. Details of the datasets are included in the supplementary material.

5.2 Results and Discussions

In this section, we study the speed-accuracy trade-off and the impact of the length model. Additional ablation experiments are included in the supplementary material.

Speed vs. Accuracy Trade-off. One of the major benefits of FIFA is that it is any-time. It provides the flexibility of choosing the number of optimization steps. The number of steps of the optimization can be a tool to trade-off speed vs. accuracy. In exact inference, we can use frame-sampling, i.e., lowering the resolution of the input features, or hypothesis pruning, i.e., beam search for speed vs. accuracy trade-off.

Figure 2 plots the speed vs. accuracy trade-off of exact inference compared to FIFA. We observe that FIFA provides a much better speed-accuracy trade-off as compared to frame-sampling for exact inference. The best performance is achieved after 50 steps with 5.9% improvement on the MoF accuracy compared to not performing any optimization (0 steps).

Impact of the Length Energy Multiplier. For the length energy, we assume that the segment lengths follow a Laplace distribution. Figure 3 shows the impact of the length energy multiplier on the weakly supervised action segmentation performance on the Breakfast dataset. The choice of this parameter depends on the dataset. While the best accuracy is achieved with a multiplier of 0.05, our approach is robust to the choice of these hyper-parameters on this dataset. We further experimented with a Gaussian length energy. However, as shown in the figure, the performance is much worse compared to the Laplace energy. This is due to the quadratic penalty that dominates the total energy, which makes the optimization biased towards the initial estimate and ignores the observation energy.

Impact of the Length Initialization. Since FIFA starts with an initial estimate for the lengths, the choice of initialization might have an impact on the performance. Table 1 shows the effect of initializing the lengths with equal values compared to using the length model of MuCon [34] for the weakly supervised action segmentation on the Breakfast dataset. As shown in the table, FIFA is more robust to initialization compared to the exact inference as the drop in performance is approximately half of the exact inference.

Inference Method Initialization		MoF
Exact	MuCon [34]	50.7
	Equal	48.8 (-1.9)
FIFA	MuCon [34]	51.3
	Equal	50.2 (-1.1)

Table 1. Impact of the length initialization for MuCon using exact inference and FIFA for weakly supervised action segmentation on the Breakfast dataset.

5.3 Comparison to State of the Art

In this section, we compare FIFA to other state-of-the-art approaches.

Weakly Supervised Action Segmentation. We apply FIFA on top of two state-of-the-art approaches for weakly supervised action segmentation namely MuCon [34] and CDFL [27] on the Breakfast dataset [17] and report the results in Table 2. FIFA applied on CDFL achieves a 12 times faster inference speed while obtaining results comparable to exact inference. FIFA applied to MuCon achieves a 5 times faster inference speed and obtains a new state-of-the-art performance on the Breakfast dataset on most of the metrics.

We also reported the inference speed of ISBA [7] and NNV [32] (since the source code of D3TW [6] is not available we could not measure its inference speed) and reported them for the sake of completeness in Table 2. ISBA has the fastest inference time as during testing it does not perform any optimization. ISBA makes framewise predictions which results in over-segmentation and low performance across all metrics.

Method	MoF	MoF-BG	IoU	IoD	Time (min)
ISBA [7]	38.4	38.4	24.2	40.6	0.01
NNV [32]	43.0	-	-	-	234
D3TW [6]	45.7	-	-	-	-
CDFL [27]	50.2	48.0	33.7	45.4	-
CDFL*	49.4	47.5	35.2	46.4	260
FIFA + CDFL*	47.9	46.3	34.7	48.0	20.4 ($\times 12.8$)
MuCon [34]	47.1	-	-	-	-
MuCon*	<u>50.7</u>	<u>50.3</u>	<u>40.9</u>	54.0	4.1
FIFA + MuCon*	51.3	50.7	41.1	<u>53.3</u>	0.8 ($\times 5.1$)

Table 2. Results for weakly supervised action segmentation on the Breakfast dataset. * indicates results obtained by running the code on our machine.

Method	F1@{10, 25, 50} Edit MoF				
BCN [36]	68.7	65.5	<u>55.0</u>	66.2	70.4
ASRF [14]	<u>74.3</u>	68.9	56.1	72.4	67.6
MS-TCN++ [28]	64.1	58.6	45.9	65.6	67.6
FIFA + MS-TCN++*	<u>74.3</u>	<u>69.0</u>	54.3	<u>77.3</u>	67.9
MS-TCN [1]	52.6	48.1	37.9	61.7	66.3
FIFA + MS-TCN*	75.5	70.2	54.8	78.5	<u>68.6</u>

Table 3. Results for fully supervised action segmentation setup on the Breakfast dataset. * indicates results obtained by running the code on our machine.

Similarly for the Hollywood extended dataset [4], we apply FIFA to MuCon [34] and report the results in Table 4. FIFA applied on MuCon achieves a 4 times faster inference speed while obtaining results comparable to MuCon with exact inference.

Fully Supervised Action Segmentation. In the fully supervised action segmentation setting, we apply FIFA on top of MS-TCN [1] and its variant MS-TCN++ [28] on the Breakfast dataset [17] and report the results in Table 3. MS-TCN and MS-TCN++ do not perform any inference at test time. This usually results in over-segmentation and low F1 and Edit scores. Applying FIFA on top of these approaches improves the F1 and Edit scores significantly. FIFA applied on top of MS-TCN achieves state-of-the-art performance on most metrics.

For the Hollywood extended dataset [4], we train MS-TCN [1] and report results comparing exact inference (EI) compared to FIFA in Table 5. We observe that MS-TCN using an inference algorithm achieves new state-of-the-art results on this dataset. FIFA is comparable or better than exact inference on this dataset.

Method	MoF-BG	IoU	Time (speedup)
ISBA [7]	34.5	12.6	-
D3TW [27]	33.6	-	-
CDFL [27]	40.6	19.5	-
MuCon [34]	41.6	-	-
MuCon*	40.1	<u>13.9</u>	53
MuCon + FIFA*	<u>41.2</u>	13.7	13 ($\times 4.1$)

Table 4. Results for weakly supervised action segmentation on the Hollywood extended dataset. Time is reported in seconds. * indicates results obtained by running the code on our machine.

Method	MoF	MoF-BG	IoU	IoD
HTK [19]	39.5	-	8.4	-
ED-TCN [7]	36.7	27.3	10.9	13.1
ISBA [7]	54.8	33.1	20.4	28.8
MSTCN [34] (+ EI)*	<u>64.9</u>	35.0	<u>22.6</u>	<u>33.2</u>
MSTCN + FIFA*	66.2	<u>34.8</u>	23.9	35.8

Table 5. Results for fully supervised action segmentation on the Hollywood extended dataset. * indicates results obtained by running the code on our machine. EI stands for Exact Inference.

5.4 Qualitative Example

A qualitative example of the FIFA optimization process is depicted in Figure 4. For further qualitative examples, failure cases, and details please refer to the supplementary material.

6 Conclusion

In this paper, we proposed FIFA a fast approximate inference procedure for action segmentation and alignment. Unlike previous methods, our approach does not rely on dynamic programming-based Viterbi decoding for inference. Instead, FIFA optimizes a differentiable energy function that can be minimized using gradient-descent which allows for a fast and accurate inference during testing. We evaluated FIFA on top of fully and weakly supervised methods trained on the Breakfast and Hollywood extended datasets. The results show that FIFA is able to achieve comparable or better performance, while being at least 5 times faster than exact inference.

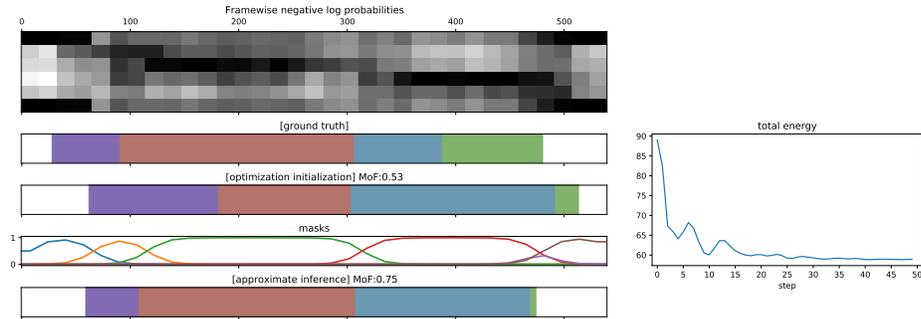


Fig. 4. Visualization of the FIFA optimization process. On the right the values of the total approximate energy is plotted. On the left, negative log probability values, ground truth segmentation, optimization initialization, the masks and the segmentation after inference are plotted.

Acknowledgement. The work has been partially funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) - GA 1927/4-2 (FOR 2535 Anticipating Human Behavior).

References

1. Abu Farha, Y., Gall, J.: Ms-tcn: Multi-stage temporal convolutional network for action segmentation. CVPR (2019)
2. Belanger, D., McCallum, A.: Structured prediction energy networks. In: ICML (2016)
3. Belanger, D., Yang, B., McCallum, A.: End-to-end learning for structured prediction energy networks. In: ICML (2017)
4. Bojanowski, P., Lajugie, R., Bach, F., Laptev, I., Ponce, J., Schmid, C., Sivic, J.: Weakly supervised action labeling in videos under ordering constraints. In: ECCV (2014)
5. Bojanowski, P., Lajugie, R., Bach, F., Laptev, I., Ponce, J., Schmid, C., Sivic, J.: Weakly supervised action labeling in videos under ordering constraints. In: ECCV (2014)
6. Chang, C., Huang, D., Sui, Y., Fei-Fei, L., Niebles, J.C.: D³TW: Discriminative differentiable dynamic time warping for weakly supervised action alignment and segmentation. CVPR (2019)
7. Ding, L., Xu, C.: Weakly-supervised action segmentation with iterative soft boundary assignment. In: CVPR (2018)
8. Duchenne, O., Laptev, I., Sivic, J., Bach, F., Ponce, J.: Automatic annotation of human actions in video. In: ICCV (2009)
9. Gatys, L.A., Ecker, A.S., Bethge, M.: A neural algorithm of artistic style. arXiv preprint arXiv:1508.06576 (2015)
10. Goodfellow, I.J., Shlens, J., Szegedy, C.: Explaining and harnessing adversarial examples. In: ICLR (2015)
11. Hoang, C.D.V., Haffari, G., Cohn, T.: Towards decoding as continuous optimization in neural machine translation. In: EMNLP (2017)

12. Huang, D.A., Fei-Fei, L., Niebles, J.C.: Connectionist temporal modeling for weakly supervised action labeling. In: ECCV (2016)
13. Huang, Y., Sugano, Y., Sato, Y.: Improving action segmentation via graph-based temporal reasoning. In: CVPR (2020)
14. Ishikawa, Y., Kasai, S., Aoki, Y., Kataoka, H.: Alleviating over-segmentation errors by detecting action boundaries. In: WACV (2021)
15. Johnson, J., Alahi, A., Fei-Fei, L.: Perceptual losses for real-time style transfer and super-resolution. In: ECCV (2016)
16. Karaman, S., Seidenari, L., Del Bimbo, A.: Fast saliency based pooling of fisher encoded dense trajectories. In: ECCV THUMOS Workshop (2014)
17. Kuehne, H., Arslan, A., Serre, T.: The language of actions: Recovering the syntax and semantics of goal-directed human activities. In: CVPR (2014)
18. Kuehne, H., Gall, J., Serre, T.: An end-to-end generative framework for video segmentation and recognition. In: WACV (2016)
19. Kuehne, H., Richard, A., Gall, J.: Weakly supervised learning of actions from transcripts. CVIU (2017)
20. Kuehne, H., Richard, A., Gall, J.: Weakly supervised learning of actions from transcripts. *Computer Vision and Image Understanding* **163**, 78–89 (2017)
21. Kuehne, H., Richard, A., Gall, J.: A Hybrid RNN-HMM approach for weakly supervised temporal action segmentation. *PAMI* **42**(04), 765–779 (2020)
22. Le, Q., Mikolov, T.: Distributed representations of sentences and documents. In: ICML (2014)
23. Lea, C., Flynn, M.D., Vidal, R., Reiter, A., Hager, G.D.: Temporal convolutional networks for action segmentation and detection. In: CVPR (2017)
24. Lea, C., Reiter, A., Vidal, R., Hager, G.D.: Segmental spatiotemporal cnns for fine-grained action segmentation. In: ECCV (2016)
25. LeCun, Y., Chopra, S., Hadsell, R., Ranzato, M., Huang, F.: A tutorial on energy-based learning. *Predicting structured data* **1** (2006)
26. Lei, P., Todorovic, S.: Temporal deformable residual networks for action segmentation in videos. In: CVPR (2018)
27. Li, J., Lei, P., Todorovic, S.: Weakly supervised energy-based learning for action segmentation. In: ICCV (2019)
28. Li, S., Abu Farha, Y., Liu, Y., Cheng, M.M., Gall, J.: MS-TCN++: Multi-stage temporal convolutional network for action segmentation. *PAMI* (2020)
29. Moltisanti, D., Fidler, S., Damen, D.: Action recognition from single timestamp supervision in untrimmed videos. In: CVPR (2019)
30. Pirsiavash, H., Ramanan, D.: Parsing videos of actions with segmental grammars. In: CVPR (2014)
31. Richard, A., Kuehne, H., Gall, J.: Weakly supervised action learning with rnn based fine-to-coarse modeling. In: CVPR (2017)
32. Richard, A., Kuehne, H., Iqbal, A., Gall, J.: Neuralnetwork-viterbi: A framework for weakly supervised video learning. In: CVPR (2018)
33. Rohrbach, M., Amin, S., Andriluka, M., Schiele, B.: A database for fine grained activity detection of cooking activities. In: CVPR (2012)
34. Souri, Y., Fayyaz, M., Minciullo, L., Francesca, G., Gall, J.: Fast Weakly Supervised Action Segmentation Using Mutual Consistency. *PAMI* (2021)
35. Vo, N.N., Bobick, A.F.: From stochastic grammar to bayes network: Probabilistic parsing of complex activity. In: CVPR (2014)
36. Wang, Z., Gao, Z., Wang, L., Li, Z., Wu, G.: Boundary-aware cascade networks for temporal action segmentation. In: ECCV (2020)