

MIT Open Access Articles

Fast online segmentation of activities from partial trajectories

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Iqbal, Tariq, et al., "Fast online segmentation of activities from partial trajectories." 2019 International Conference on Robotics and Automation (ICRA), May 20-24, 2019, Montreal, QC, IEEE, 2019: pp. 5019-25 doi 10.1109/ICRA.2019.8794054 ©2019 Author(s)

As Published: 10.1109/ICRA.2019.8794054

Publisher: IEEE

Persistent URL: <https://hdl.handle.net/1721.1/125900>

Version: Author's final manuscript: final author's manuscript post peer review, without publisher's formatting or copy editing

Terms of use: Creative Commons Attribution-Noncommercial-Share Alike



Fast Online Segmentation of Activities from Partial Trajectories

Tariq Iqbal¹, Shen Li¹, Christopher Fourie¹, Bradley Hayes², and Julie A. Shah¹

Abstract—Augmenting a robot with the capacity to understand the activities of the people it collaborates with in order to then label and segment those activities allows the robot to generate an efficient and safe plan for performing its own actions. In this work, we introduce an online activity segmentation algorithm that can detect activity segments by processing a partial trajectory. We model the transitions through activities as a hidden Markov model, which runs online by implementing an efficient particle-filtering approach to infer the maximum a posteriori estimate of the activity sequence. This process is complemented by an online search process to refine activity segments using task model information about the partial order of activities. We evaluated our algorithm by comparing its performance to two state-of-the-art activity segmentation algorithms on three human activity datasets. The proposed algorithm improved activity segmentation accuracy across all three datasets compared with the other two approaches, with a range from 11.3% to 65.5%, and could accurately recognize an activity through observation alone for 31.6% of the initial trajectory of that activity, on average. We also implemented the algorithm onto an industrial mobile robot during an automotive assembly task in which the robot tracked a human worker’s progress and provided the worker with the correct materials at the appropriate time.

I. INTRODUCTION

Robots currently have the capacity to help people in several fields, including health care, assisted living, and manufacturing and factory settings. In many of these scenarios, robots must share physical space and actively collaborate with humans [1–4]. The performance of many of these human-robot teams depends upon how fluently all team members can jointly perform their tasks [5–7]. In order to successfully act within a group, people must be able to predict the intentions of other group members and use that knowledge to determine when, where, and how to act for the team’s benefit [8]. In human-robot interaction scenarios, a robot similarly requires the ability to precisely identify and monitor other members’ actions so that it can predict future actions and adapt its own plans accordingly [9, 10]. In particular, a robot requires the ability to segment others’ activities online by detecting the start time of each activity and distinguishing it from the end time of the previous activity. This capacity is crucial to efficient and safe human-robot interactions within factory environments, where humans and robots often work in close physical proximity to one another [11].

Researchers across many fields have attempted to address this concern, which is defined as the *online activity*

¹The authors are with the Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, MA, USA. {tiqbal, shenli, ckfourie}@mit.edu, julie.a.shah@csail.mit.edu. ²The author is with the Dept. of Computer Science, University of Colorado Boulder, Boulder, CO, USA. bradley.hayes@colorado.edu.



Fig. 1: Three activities similar to the activities from the Auto-DA dataset: A) moving to the dashboard, B) collecting the speedometer unit, C) placing the speedometer unit onto the dashboard.

segmentation problem. Prior techniques [12–15] are well-suited for data segmentation via post-processing but are not applicable to human-robot interaction scenarios, because either the algorithms become computationally expensive to run online, or they cannot function adequately with only a partial observance of data.

In this paper, we introduce FOSAPT (Fast Online Segmentation of Activities from Partial Trajectories), an online activity-segmentation algorithm able to identify and label activity segments by processing just a part of the full activity trajectory. We model the transitions through the activity classes as a hidden Markov model (HMM) in a manner similar to Fearnhead and Liu’s approach [16]. FOSAPT runs efficiently in an online setting by inferring the maximum a posteriori estimate of the sequence of activity classes, using particle filtering to make the inference tractable. However, FOSAPT’s particle-filter approach is complemented by an online search process to refine activity segments via task model information about the partial order of activities, along with predictive models of the timings of future activities.

We evaluated FOSAPT by assessing its performance on three human activity datasets (UTKinect [17], Static-Reach [18], and Auto-DA) compared with two state-of-the-art activity segmentation algorithms from prior literature: the online change point detection (CPD) algorithm [16] and the transition state clustering algorithm [12]. One of the datasets, Auto-DA, is derived from a real-world manufacturing task involved in automotive final assembly, in which 12 manufacturing associates performed variations of the task (see Fig. 1).

In our empirical evaluation, FOSAPT accurately segmented 72.2% of the activity segments, on average, across all four variants of the task orders in the Auto-DA dataset. We also observed that FOSAPT could accurately recognize and label an activity by observing only 31.6% of that activity’s initial trajectory, on average. FOSAPT achieved an improvement to accuracy ranging from 11.3% to 65.5% compared with the other two approaches across all three datasets. Finally, we applied our algorithm to a demonstration in which a collaborative robot tracked a human worker’s progress during a dashboard assembly task and provided the worker with the correct materials for assembly at the appropriate time.

II. RELATED WORK

Researchers across many fields have attempted to address the problem of activity segmentation, with two primary approaches having been explored thus far: supervised and unsupervised learning of segmentation models. Many supervised approaches involve two primary steps for activity segmentation: extracting spatiotemporal features from the data source, then training temporal models to identify segments [19–22]. For example, Lea et al. [19] developed temporal convolutional networks (TCNs), a class of time-series model, that uses a hierarchy of temporal convolutions to perform fine-grained action segmentation. These methods work well when full trajectory information is available during the segmentation process; however, they are not suitable for online segmentation scenarios in which segmentation must be performed on partial data, which is the focus of our interest.

On the other hand, the unsupervised segmentation approaches attempt to discover underlying structures within the data to cluster similar trajectories together [12, 23–25]. For example, Krishnan et al. [12] introduced the transition state clustering (TSC) method, an unsupervised segmentation approach that clusters transition states together from a set of demonstrations. TSC assumes that the low-level dynamics of a segment are noisy, but that high-level dynamics follow a consistent, spatially and temporally correlated partial order of events across demonstrations. The models used in many unsupervised approaches assume particular structures within the prior probabilities and trajectory data, and an unsupervised method designed for one dataset often fails to perform satisfactorily on another dataset with a different underlying structure.

Another prominent segmentation approach involves statistical model-based changepoint detection algorithms [14, 16, 26–28]. For example, Fox et al. [14] developed the Beta Process Autoregressive HMM (BP-AR-HMM), a Bayesian, non-parametric approach to jointly modeling multiple related trajectory time-series. In the same vein, Fearnhead and Liu [16, 27] developed a statistical approach to online changepoint detection problems by introducing a sampling method similar to particle filters to reduce computational cost. They assumed that if an observation sequence and a set of candidate models are given, then the observation sequence is generated from specific underlying models, and the points at which the underlying models change are detected as changepoints. Konidaris et al. [28] built upon this idea and implemented a model-based changepoint detection process for constructing skill trees to acquire skills from human demonstrations.

Although these methods work for activity segmentation in various scenarios, they either become computationally expensive to run online or do not perform adequately with only partial observance of data. In order to address these problems, we incorporate a supervised learning model with a statistical changepoint detection algorithm and a particle filtering-based sampling approach to performing online human activity segmentation onto a robot. This algorithm is able to function with only partial observance of trajectory data.

III. METHODOLOGY

In this section, we introduce the FOSAPT (Fast Online Segmentation of Activities from Partial Trajectories) algorithm. This segmentation algorithm takes a partial trajectory as input and can recognize and track activity segments as it progresses in real time. FOSAPT is also capable of further refining the start and the end times of an activity accurately as that activity ends. This approach is both computationally inexpensive and suitable for real-time implementation onto robots.

A. Problem Formulation

Given a set of trajectory spanning time $t = 1$ to T with corresponding trajectory frames $F = (f_1, f_2, \dots, f_T)$ and a set of activity labels $A = (a_1, a_2, \dots, a_m)$, the activity segmentation problem is defined as identifying the mapping of trajectory frames to activity labels. In this work, we are interested in online activity segmentation with only partial observance of trajectory data.

We define the times at which an activity label changes in the trajectory frames as denoted by changepoints, τ . Consider that the trajectory data consists of n activities; we can then denote $(0 < \tau_1 < \tau_2 < \dots < \tau_{n-1} < T)$ and $\tau_0 = 0$ and $\tau_n = T$ as the changepoints in trajectory frames.

The trajectory frames between two changepoints are denoted as a segment, s . We define a segment of frames as $s(i, j) = (f_i, \dots, f_j)$, where $f \in F$ and $i = \tau_p$ and $j = \tau_{p+1}$. Thus, given a set of activities A , activity segmentation produces a set S that contains an allocation of non-overlapping frame intervals in F with labels drawn from A . Thus, $s(i, j, a) \forall s \in S, i, j \in T$ such that $i < j, a \in A$.

Online partial trajectory activity segmentation assumes a partial set of trajectory frames $s(i, k) = (f_i, \dots, f_k)$ as input, where $i = \tau_p, k \leq j = \tau_{p+1}$. The output is $s(i, k, a)$ such that $s(i, k, a) \subseteq s(i, j, a), (i, k, j) \in T, (i < k \leq j),$ and $a \in A$.

B. Approach Overview

FOSAPT relies upon a small portion of data as it arrives and checks the likelihood of the label of that segment via a set of trained activity classifiers. It then incrementally computes the likelihood value of a larger segment from smaller segment likelihoods (instead of calculating the likelihood of the larger segment again). In FOSAPT, transitions through the activity classes are modeled as a hidden Markov model. It runs efficiently in an online setting by inferring the maximum a posteriori estimate of the sequence of activity classes via a particle-filtering approach. This method is complemented by an online search process to refine activity segments using task model information about the partial order of activities, as well as predictive models of the timings of future activities. We present the algorithm in Algorithm 1.

C. Activity Likelihood Calculation from Partial Trajectories

FOSAPT utilizes an activity classification algorithm to measure the activity likelihood of a set of trajectory frames, described as the *likelihood* function (Algo 1 - Line 5). In this implementation, we utilize RAPTOR [29], a real-time, state-of-the-art activity classifiers; however, FOSAPT is agnostic to the choice of the activity classification algorithm.

Algorithm 1 FOSAPT

Input: Partial trajectory $\langle \Delta F, j, t \rangle$, Activity set $\langle A \rangle$, Activity classifiers $\langle C \rangle$, Number of sub-classes $\langle B \rangle$, Bin lengths $\langle \beta \rangle$, Particles $\langle P \rangle$, Previously detected segments $\langle S \rangle$, Future activity list $\langle FAL \rangle$
Output: Activity segments $\langle S \rangle$, Particles $\langle P \rangle$, Bin lengths $\langle \beta \rangle$

```

1:  $pList \leftarrow \phi$   $\triangleright$  temporary particle list
2: for  $a \in FAL$  do
3:   if  $(|\Delta F| \geq \beta(a))$  then
4:     for  $b \leftarrow 1 : B$  do
5:        $l \leftarrow \text{likelihood}(C_a^b, \Delta F)$   $\triangleright$  classifier fitness
6:        $pr \leftarrow \text{prior}(a, FAL)$   $\triangleright$  task structure
7:        $pmp \leftarrow \text{prev\_MAP}(P, j, t, A)$   $\triangleright$  find MAP particle ends at  $j$ 
8:        $P_t(j, a) \leftarrow (1 - G_a(t - j - 1)) * l * pr * pmp.MAP$ 
9:        $MAP \leftarrow P_t(j, a) * g_a(t - j) / (1 - G_a(t - j - 1))$ 
10:       $np \leftarrow \text{Particle}(a, b, j, t, l, P_t(j, a), MAP, pmp)$ 
11:       $pList = pList + np$ 
12: for  $p \in pList$  do
13:    $P \leftarrow \text{merge\_or\_insert}(p, P)$   $\triangleright$  merge with other particles, or insert in  $P$ 
14:    $P \leftarrow \text{resample}(P)$   $\triangleright$  resample particles
15:    $S \leftarrow \text{find\_all\_segments}(P, S)$   $\triangleright$  find all activity segments
16:    $\beta \leftarrow \text{dynamic\_bin\_length\_adjustments}(\beta, S)$   $\triangleright$  adjust bin lengths from evidence
17:    $\text{refine\_activity\_segments}(S)$   $\triangleright$  this step runs in parallel
18: return  $S, P, \beta$ 

```

We first train a set of activity classifiers $C = (c_{a_1}, c_{a_2}, \dots, c_{a_m})$ for each $a_i \in A$, where $m = |A|$, on a corpus of recorded trajectory frames, such that activity classifier c_{a_i} is trained for activity a_i . During the training phase, we train RAPTOR with an ensemble of sub-classifiers using a part of the full trajectory frames, and combine those temporally to generate a full activity classifier, similar to the idea proposed by Hayes and Shah [29]. We divide activity a_i into B sub-classes. We then train $(c_{a_i}^1, c_{a_i}^2, \dots, c_{a_i}^B)$ sub-classifiers separately and combine them temporally to generate classifier c_{a_i} for activity a_i . A training process such as this permits the algorithm to utilize sub-classifiers directly (e.g., check the likelihood of trajectory frames (f_i, \dots, f_k) with $c_{a_i}^b$, where $1 \leq b \leq B$ for activity a_i).

From the training data, FOSAPT models the marginal probability of each activity length with a probability mass function $g(\cdot)$. Thus, we can define $P(\tau_k - \tau_{(k-1)} = d) = g_{a_i}(d)$, where $g(\cdot)$ is a discrete distribution on the length of activity a_i . The corresponding cumulative distribution function is $G_{a_i}(d) = \sum_{f=1}^d g_{a_i}(f)$ [16]. Each activity classifier is a temporal ensemble of B sub-classifiers, and FOSAPT models the bin lengths as β from the training data, where $\beta(a)$ denotes the bin length of activity a . We explain how to dynamically adjust this value in Section III-J.

During the testing phase, FOSAPT utilizes these sub-classifiers to measure the likelihood of a partial set of trajectory frames being an activity. If ΔF denotes a partial set of trajectory frames (f_{j+1}, \dots, f_t) , then the algorithm checks whether the length of the trajectory frames $(|\Delta F|)$ is greater than or equal to the bin length of that activity, before performing the likelihood computation (Line 3). FOSAPT then tests the likelihood of ΔF being an activity a with each sub-classifier C_a^b , where $1 \leq b \leq B$ (Lines 4 and 5).

D. Task Structure Modeling

For a set of trajectory frames (ΔF) , if the algorithm needs to check all possible combinations of activity classes ($\forall c_a$ for $a \in A$), then the process would be computationally expensive. Thus, FOSAPT leverages prior knowledge of the task's sequence of activities to reduce the number of classifiers

on which the trajectory frames are tested. Thus, a classifier is only used to evaluate a trajectory when a task is supposed to occur at that moment or at a time in the near future.

The algorithm constructs a hierarchical task network similar to the clique/chain hierarchical task network (CC-HTN) proposed by Hayes et al. [30]. FOSAPT builds this network using the high-level task sequences of activities from the training demonstrations, and generates a *future_activity_list* (FAL) incorporating the possible activities that could happen at a given moment of time from this task network, including the current activity. Thus, FOSAPT only tests the trajectory frames for activities that are present in the FAL (Line 2), reducing the number of expensive *likelihood* computations.

E. Activity Prior Computation

The algorithm computes a *prior* probability ($pr(a)$) for an activity (a), described as *prior*(a, FAL) (Line 6), utilizing the FAL . This value represents the probability of an activity occurring at that time. One can learn any prior probability distribution of an activity from the training data. In our case, the algorithm assumes that all the activities in the FAL are equally likely; thus, the algorithm is initialized with a uniform prior.

F. Activity Transition Modeling

Taking inspiration from the statistical changepoint detection method developed by Fearnhead and Liu [16, 27], FOSAPT can utilize these values to model the activity transitions as a hidden Markov model. Here, the observed state at time t is the trajectory frame, $\{f_t\}$, and the hidden state is the activity label, a_t . The probability of a set of trajectory frames that starts at time $(j+1)$ and ends at t ($\Delta F = f_{j+1}, \dots, f_t$) being activity a_t can be modeled as the product of the likelihood of ΔF being activity a_t (l , measured in Line 5) and the probability of the segment lasting for $(t - j - 1)$ time steps. Thus, this probability can be defined as $P(\Delta F|a_t) = l * (1 - G_{a_t}(t - j - 1))$. Similarly, the transition probability of the activities can be defined as $T(a_j, a_t) = g_{a_t}(t - j - 1) * pr(a_t)$, where the transition from activity a_j at time j to activity a_t occurs at time t , and $pr(a_t)$ represents the prior probability of activity a_t (Line 6) [16, 28].

FOSAPT can now compute the maximum likelihood sequence for the activities (hidden states) given their transition probability and the trajectory frames (observation). This enables the algorithm to use an online Viterbi algorithm to calculate the MAP estimate of the activity changepoint positions and the orders. Thus, we compute the probability of an activity a starting at time $(j + 1)$ and continuing at time t ($P_t(j, a)$) as follows:

$$P_t(j, a) = (1 - G_a(t - j - 1)) * l * pr(a) * P_j^{MAP} \quad (1)$$

$$P_t^{MAP} = \max_{j, a} \frac{P_t(j, a) g_a(t - j)}{(1 - G_a(t - j - 1))} \quad (2)$$

A particle filter is employed to efficiently keep the computation tractable [16], where each particle represents a segment with activity a that starts at time $(j + 1)$ and ends at time t . Each particle stores the activity (a), the start and the end bin numbers (b), the start time and the end of that activity segment (j and t), the activity likelihood value (l), $P_t(j, a)$,




Fig. 2: A depiction of the particle generation and merging process. Left) Particles are generated by computing the sub-classifier likelihoods from ΔF_1 frames. Center) New particles are generated from ΔF_2 frames (green) and merged with the previously generated particles (orange). Right) Merged, new, and old particles are in the particle list (new particles from ΔF_2 in green, merged particles in blue, and old particles from ΔF_1 in orange).

the MAP value, and the maximum likelihood sequence. FOSAPT first computes the MAP particle that ends at time j from previously generated particles (Line 7), and then computes $P_t(j, a)$ and the MAP value according to Eq. 1 and Eq. 2 (Lines 8-9). It then generates a new particle and stores it in a temporary particle list (Lines 10-11).

G. Incremental $P_t(j, a)$ Computation Through Merging

Consider a situation in which a_i is currently executing. Given a small set of trajectory frames, $\Delta F_1 \sim (f_{i_1}, \dots, f_{j_1})$, FOSAPT does not have any information regarding how far the activity has already progressed, which would enable it to check a specific sub-classifier. Thus, it checks the likelihood of the segment ΔF_1 by testing against all the sub-classifiers ($c_{a_i}^b$) of that activity where $b = 1, \dots, B$ (Line 4). After measuring the likelihood values for all sub-classifiers, FOSAPT generates B number of particles, and stores all particles in a particle list, P (Line 13, *insert* operation). Suppose these particles are called $(c_{a_i}^1(\Delta F_1), c_{a_i}^2(\Delta F_1), \dots, c_{a_i}^B(\Delta F_1))$.

FOSAPT then waits for another set of trajectory frames – for example, $\Delta F_2 \sim (f_{i_2}, \dots, f_{j_2})$. Following a similar process, the algorithm then generates another B particles, with each particle storing the likelihood of segment ΔF_2 tested against all sub-classifiers; these particles are called $(c_{a_i}^1(\Delta F_2), c_{a_i}^2(\Delta F_2), \dots, c_{a_i}^B(\Delta F_2))$.

Next, the algorithm performs a *merge* operation. The idea behind this merge is that the algorithm computes the likelihood of a larger segment from already-computed likelihoods of smaller segments. This operation is computationally inexpensive, but enables the algorithm to efficiently compute the likelihood of larger segments while concurrently maintaining multiple possible hypotheses.

Particles must be temporally adjacent in order for the algorithm to merge them. Two particles are considered temporally adjacent when they contain the likelihood values computed from two adjacent sub-classifiers of the same activity. For example, the algorithm would merge particles $c_{a_i}^1(\Delta F_1)$ and $c_{a_i}^2(\Delta F_2)$, as the likelihood value of particle $c_{a_i}^1(\Delta F_1)$ is computed from sub-classifier 1 ($c_{a_i}^1$) and the likelihood value of particle $c_{a_i}^2(\Delta F_2)$ is computed from sub-classifier 2 ($c_{a_i}^2$). As such, during this process, the algorithm merges $c_{a_i}^2(\Delta F_2)$ with $c_{a_i}^1(\Delta F_1)$, and the $c_{a_i}^1(\Delta F_1)$ particle becomes $c_{a_i}^{(1-2)}(\Delta F_{(1-2)})$, while $c_{a_i}^2(\Delta F_2)$ does not change.

To compute the $P_t(j, a)$ of this merged particle

($c_{a_i}^{(1-2)}(\Delta F_{(1-2)})$), FOSAPT utilizes the computed likelihood values of the $c_{a_i}^1(\Delta F_1)$ and $c_{a_i}^2(\Delta F_2)$ particles. As the trajectory segments are temporally adjacent but independent, the activity likelihood values of the merged trajectory segment ($\Delta F_1 + \Delta F_2$) are computed by taking a product of the activity likelihood of both segments and then normalizing it over the number of segments (Line 5). Similar to RAPTOR, FOSAPT performs a max-pooling operation on the likelihood values before performing this computation. The algorithm updates the values of other statistics by following the steps presented in Lines 6-9 for the merged segment. (We present this merging step in Fig. 2.) After merging the eligible particles, the algorithm updates the particle list, P .

H. Resampling

To keep the computation tractable, when the number of particles ($|P|$) reaches a predefined threshold (R_N), FOSAPT applies a resampling algorithm to reduce the number of particles to a fixed value (R_M) (Line 14). First, the algorithm removes the particles with a starting sub-classifier number higher than a threshold value, as these particles represent the final portion of an activity and are not likely to produce a larger segment to represent a major part of the activity in the future. Second, the algorithm removes any segments that are not updated for a set time period. In our implementation, FOSAPT measured the longest time each activity takes from the training data, summed this value with the square root of the standard deviation value of that activity time, and set the result as the maximum allowable time for a particle of that activity to be alive. Finally, after applying both of the previous filtering methods, if the total number of particles is higher than R_M , FOSAPT prunes $m\%$ of the particles with the lowest likelihood.

I. Compute Activity Segments

After completing the particle generation process, the algorithm's goal is to identify the activity sequence from the observed trajectory up to that time step. To do so, FOSAPT first finds the particle within the particle list (P) with the highest $P_t(j, a)$ value at time t . We impose additional constraints that a particle must be longer than a predefined number of bin lengths (β) and the likelihood value of that particle must be greater than a likelihood threshold value (ϕ) before it can be selected as a candidate for an activity. After finding this particle, the algorithm adds this particle to a *max_particle_path*. All the activity segments are computed by backtracking the *max_particle_path* (Line 15).

J. Activity Time Prediction and Dynamic β Adjustment

FOSAPT updates its predictions about activity timing to improve segmentation performance. If a test activity trajectory is longer or shorter than the training trajectories, the bin length (β) of that activity must be adjusted accordingly to compute accurate likelihood values from the sub-classifiers. To address this challenge, we designed an algorithm that generates predictions about the timing of future activities based on observed timings of previous activities and adjusts the bin length (β) according to these predictions.

A Gaussian Mixture Model (GMM) is used to model temporal predictions of an activity time. To facilitate prediction based on learned data, we use the conditional decomposition of the model. This allows the algorithm to incorporate accumulated evidence as it is obtained from the activity segmentation. The conditional of the GMM can be calculated as follows: $p(\mathbf{x}_s|\mathbf{x}_d) = \sum_k w_k \frac{p_k(\mathbf{x}_d)}{p(\mathbf{x}_d)} \cdot p_k(\mathbf{x}_s|\mathbf{x}_d)$. Here, \mathbf{x}_s represents the variables over which inference is to be performed, and \mathbf{x}_d represents the variables for which evidence exists. This equation illustrates that the conditional of a sum is the weighted sum of the individual conditionals of the mode, with the weight defined by a ratio of the marginals. FOSAPT trains these models during the training phase.

As more information is obtained during the testing phase, the model can be updated to form a new distribution with different means and variances. The conventional method for forming a prediction would be to use the maximum of the conditional distribution. However, to hasten inference, FOSAPT incorporates the mean of the highest weighted mode in lieu of the true maximum, which is interpreted as the prediction.

After computing the timing of the current activity and predicting the timing of future activities, FOSAPT adjusts the bin length (β) of each activity, while the number of sub-classifiers (B) remains the same. This dynamic bin length adjustment process (Line 16) makes the algorithm robust to any deviations from the trained models with regard to activity timing.

K. Refine Activity Segments

As activity segmentation progresses, FOSAPT performs another concurrent search within previously detected activity segments to refine its assessment of the likely start and end time points. This search helps to determine the accurate transition points to nearby activities.

To search for an accurate bound, the algorithm moves the beginning and ending of a previously detect segment forward and backward to identify the part with the highest likelihood value. As this segmentation search is only employed on completed activity segments, it does not interfere with the current activity segmentation approach; this step runs in parallel with the current activity segmentation steps. To reduce computation, this step is only executed on an activity segment when that segment does not change in the *activity segment list*, S , for a predefined amount of time. Were the activity sequence to change in the future, the refinement process would be executed again on S (Line 17). The overall time complexity of FOSAPT is $O(|FAL| * B * |\Delta F| + |P|)$.

IV. EXPERIMENTAL VALIDATION

We evaluated FOSAPT’s performance by comparing it against two activity segmentation methods from prior literature: the Change Point Detection (CPD) algorithm developed by Fearnhead and Liu [16], which uses a similar statistical approach to FOSAPT for tracking activities; and the Transition State Clustering (TSC) algorithm by Krishnan et al. [12], a state-of-the-art activity-clustering algorithm. In our implementation, the CPD algorithm used RAPTOR to determine the likelihood value of a segment, which helped us to make a fair comparison between CPD’s and FOSAPT’s performances.

We also evaluated FOSAPT’s performance using three activity datasets: a motion-capture dataset of the automotive dashboard assembly process of real industry associates, a publicly available activity dataset (UTKinect [17]), and reaching behavior during a manufacturing task (Static-Reach [18]).

In collaboration with an industry partner, we designed a test scenario for the assembly of a car dashboard (“Auto-DA” dataset). This dataset included a total of eight activities. We set up a testbed within our industry partner’s automotive assembly factory, where a total of 12 factory associates (ten males and two females) participated in a total of four variations of the assembly task. Each task sequence was 54.4 seconds long, on average. The positions and orientations of a total of seven objects (both hands, head, dashboard, speedometer, navigation unit, and a scanner gun) were tracked using a VICON motion capture system at a frame rate of 30Hz.

We combined similar activities from the UTKinect dataset [17] into the following five activity classes performed in sequence: walk, sit, pick up and transport object, shake object, and hand gestures. We removed undefined activities between labeled activities from each trial, merged all the activities sequentially for evaluation purposes, and sampled at 15Hz. One trial was not included in the data, as it did not contain all the activities on the list. The dataset contains 20 skeleton joint positions of a person tracked with a Kinect sensor.

The Static-Reach dataset was recorded via a PhaseSpace motion capture system during a human-robot collaborative task [18]. Each trial consists of 16 sequential human activities. Similar to the UTKinect dataset, we merged the sequential activities for evaluation purposes and sampled at 120Hz.

V. RESULTS

A. Evaluation metrics

We first measured the intersection-over-union (*IoU*) scores of the algorithm, following a process similar to that used in prior work ([13] and [12]). For example, if the algorithm segments an activity from a trajectory as s (representing the time duration from a starting time point to an ending point) and the ground-truth activity segment is GT , then the *IoU* is measured as $IoU = (s \cap GT)/(s \cup GT)$.

Activity segments were manually annotated for our Auto-DA dataset. We then computed segment accuracies by following an approach similar to those taken by Wu et al. [13] and Krishnan et al. [12]. We considered a segment to be detected accurately if the *IoU* value of that segment was higher than a threshold (δ). As we used real-world datasets, in keeping with Wu et al. [13], we set $\delta = 0.4$.

To evaluate the accuracy of FOSAPT and the CPD algorithm, we performed leave-one-out cross-validation across the trials for each dataset. As TSC is an unsupervised segmentation algorithm, we provided all trajectories of each dataset as input, and report the mean accuracy of five runs.

B. Accuracy of FOSAPT

We measured FOSAPT’s activity-segmentation accuracy for all task types within the Auto-DA dataset, and present the results in Table I. The results indicate that FOSAPT detected

TABLE I: Activity segmentation accuracy (%) of the FOSAPT algorithm and the percentage of initial frames of the whole trajectory it observed to accurately detect an activity applied to the Auto-DA dataset

	Segmentation accuracy (%)	Initial frames (%) to detect
<i>mv_dash</i>	78.0	14.3
<i>mv_meter</i>	76.0	23.4
<i>col_meter</i>	90.0	31.4
<i>pl_meter</i>	66.0	37.4
<i>mv_nav</i>	68.0	47.8
<i>col_nav</i>	82.0	32.5
<i>pl_nav</i>	68.0	50.1
<i>exit</i>	50.0	14.8
Average	72.2	31.6

segments with an average accuracy of 72.2% for the Auto-DA dataset, with variation from 50.0% to 90.0% depending upon the task types. For this dataset, we utilized a sequential task structure and set $|FAL| = 3$, $R_N = 120$, $R_M = 100$, $\phi = -5.0$, and $B = 15$ with 50% overlap between adjacent bins.

As the CPD algorithm had to generate particles for all activity types, the segmentation process became very slow, and it was unable to finish within 10 minutes (each task was 54.4 sec long, on average). Therefore, we incorporated the *FAL* (in Section III-D) and ϕ (in Section III-I) to keep computation tractable. We set $|FAL| = 3$ and $\phi = -5.0$, identical to the setting for FOSAPT. In keeping with the parameters used in Krishnan et al. [12], we set the parameters of TSC for all datasets, and followed a similar approach for activity label generation for the segments.

We report segmentation accuracy across all three datasets for FOSAPT, the changepoint detection (CPD) algorithm, and the transition state clustering algorithm (TSC) in Table II. The results suggest that FOSAPT was more accurate than either CPD or TSC across all three datasets, and that FOSAPT achieved an improvement in accuracy ranging from 11.3% to 65.5% vs. CPD and TSC.

We also present the number of missing activity segments reported by the algorithms in Table II. The number of missing segments was calculated by counting the number of segments from the ground-truth data that the algorithms did not report. Our findings suggest that FOSAPT failed to detect segments in fewer cases than either CPD or TSC in both the Auto-DA and UTKinect datasets (5.8% and 0.0%, respectively). CPD yielded fewer missing segments (3.8%) than FOSAPT (4.4%) for the Static-Reach dataset; however, CPD also demonstrated a lower accuracy within that dataset (20.0% accuracy compared to 88.8% with FOSAPT).

C. Performance with partial observance of trajectory data

FOSAPT is capable of detecting activities by observing just a partial trajectory of the full demonstration. We measured how long FOSAPT took to report a ground-truth activity label after the start of the activity. We then subtracted the ground-truth start time of each activity from the time when FOSAPT first reported that activity, and normalized it using the total duration of that ground-truth activity segment. Here, we only considered activity segments that were accurately segmented by FOSAPT. The results, presented in Table I, indicate that FOSAPT can segment out activities just by observing 31.6% of the full trajectory of an activity (on average).

TABLE II: Accuracy (%) and Missing Segments (%) of each algorithm across three datasets

	Auto-DA		UTKinect		Static-Reach	
	Acc.	Miss.	Acc.	Miss.	Acc.	Miss.
FOSAPT	72.2	5.8	70.5	0.0	88.8	4.4
CPD	42.9	35.0	43.4	2.1	20.0	3.8
TSC	6.7	72.0	59.2	17.5	69.6	16.0

D. Robot Demonstration

We applied FOSAPT in an industrial setting via a robot demonstration. We emulated a similar dashboard assembly scenario to the Auto-DA dataset, where, a collaborative robot tracked the human worker’s progress through the task using FOSAPT and provided the person with the right materials for assembly at the appropriate time. A video of the demonstration is available here: <http://tiny.cc/FOSAPT>.

VI. DISCUSSION AND CONCLUSION

FOSAPT outperformed all other evaluated baselines across all datasets and metrics tested. It performed best on the Static-Reach dataset (88.8% accuracy), as most of the demonstrations of this dataset followed similar paths during activity executions and exhibited relatively less jerky motion. Thus, the computed likelihood values were less ambiguous between consecutive activities, aiding in the accurate identification of activity segments. On the other hand, in the case of the UTKinect dataset (70.5% accuracy), various people performed similar actions in different ways, which could have contributed to ambiguous likelihood values for the changepoint positions and resulted in less-accurate performance by FOSAPT.

TSC demonstrated reasonably high accuracy for the UTKinect and Static-Reach datasets (59.2% and 69.6%, respectively); however, it only achieved 6.7% accuracy on Auto-DA. In the UTKinect and Static-Reach datasets, only the activity segments were temporally combined; thus, there was a substantial change to the trajectories near the activity changepoints. As TSC clusters similar data patterns, this change might contribute to better activity segmentation. However, this was not the case for the Auto-DA dataset, as it contains continuous trajectory frames for each demonstration, and TSC failed to find appropriate changepoints across the demonstrations, as each person might perform the same activity differently in space and time.

On the other hand, CPD demonstrated segmentation accuracy of 42.9% for the Auto-DA dataset and 43.4% for the UTKinect dataset, but of only 20.0% on the Static-Reach dataset. As CPD is an online algorithm, the frame rate contributed to its performance: the Static-Reach dataset had a higher frame rate (120 Hz) than the other datasets (30 Hz for Auto-DA and 15 Hz for UTKinect); thus, CPD had to generate more particles with a small degree of variation on the data, which could result in many inaccurate segment detections.

In this work, we presented FOSAPT, an online activity segmentation algorithm capable of accurately identifying and labeling activity segments, which demonstrated improved segmentation accuracy compared with two state-of-the-art segmentation algorithms, and was able to segment activities only by processing a part of the full activity trajectory (31.6% of the initial trajectory) while running online.

REFERENCES

- [1] L. Riek, "Healthcare robotics," *Communications of the ACM*, 2017.
- [2] M. J. Matarić, "Socially assistive robotics: Human augmentation versus automation," *Science Robotics*, 2017.
- [3] V. V. Unhelkar, H. C. Siu, and J. A. Shah, "Comparative performance of human and mobile robotic assistants in collaborative fetch-and-deliver tasks," in *HRI*, 2014.
- [4] T. Iqbal, S. Rack, and L. D. Riek, "Movement coordination in human-robot teams: A dynamical systems approach," *IEEE Transactions on Robotics*, vol. 32, no. 4, pp. 909–919, 2016.
- [5] T. Iqbal and L. D. Riek, "A Method for Automatic Detection of Psychomotor Entrainment," *IEEE Transactions on Affective Computing*, 2015.
- [6] S. Lemaignan, M. Warnier, E. A. Sisbot, A. Clodic, and R. Alami, "Artificial cognition for social human-robot interaction: An implementation," *Artificial Intelligence*, 2017.
- [7] T. Iqbal and L. D. Riek, "Human robot teaming: Approaches from joint action and dynamical systems," *Humanoid Robotics: A Reference, Springer Handbook of Humanoid Robotics*, 2017.
- [8] N. Sebanz, H. Bekkering, and G. Knoblich, "Joint action: bodies and minds moving together," *Trends Cogn Sci*, 2006.
- [9] S. Nikolaidis, D. Hsu, and S. Srinivasa, "Human-robot mutual adaptation in collaborative tasks: Models and experiments," *The International Journal of Robotics Research*, vol. 36, no. 5-7, pp. 618–634, 2017.
- [10] T. Iqbal and L. D. Riek, "Coordination dynamics in multi-human multi-robot teams," *IEEE Robotics and Automation Letters (RA-L)*, 2017.
- [11] P. A. Lasota, T. Fong, and J. A. Shah, "A survey of methods for safe human-robot interaction," *Foundations and Trends in Robotics*, 2017.
- [12] S. Krishnan, A. Garg, S. Patil, C. Lea, G. Hager, P. Abbeel, and K. Goldberg, "Transition state clustering: Unsupervised surgical trajectory segmentation for robot learning," *The International Journal of Robotics Research*, vol. 36, no. 13-14, pp. 1595–1618, 2017.
- [13] C. Wu, J. Zhang, O. Sener, B. Selman, S. Savarese, and A. Saxena, "Watch-n-patch: unsupervised learning of actions and relations," *IEEE transactions on pattern analysis and machine intelligence*, vol. 40, no. 2, pp. 467–481, 2018.
- [14] E. B. Fox, M. C. Hughes, E. B. Sudderth, M. I. Jordan, et al., "Joint modeling of multiple time series via the beta process with application to motion capture segmentation," *The Annals of Applied Statistics*, vol. 8, no. 3, pp. 1281–1313, 2014.
- [15] P. Bojanowski, R. Lajugie, F. Bach, I. Laptev, J. Ponce, C. Schmid, and J. Sivic, "Weakly supervised action labeling in videos under ordering constraints," in *European Conference on Computer Vision*. Springer, 2014, pp. 628–643.
- [16] P. Fearnhead and Z. Liu, "Efficient bayesian analysis of multiple changepoint models with dependence across segments," *Statistics and Computing*, 2011.
- [17] L. Xia, C. Chen, and J. Aggarwal, "View invariant human action recognition using histograms of 3D joints," in *IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, 2012, pp. 20–27.
- [18] P. A. Lasota and J. A. Shah, "Analyzing the effects of human-aware motion planning on close-proximity human-robot collaboration," *Human factors*, 2015.
- [19] C. Lea, M. D. Flynn, R. Vidal, A. Reiter, and G. D. Hager, "Temporal convolutional networks for action segmentation and detection," in *International Conference on Computer Vision (ICCV)*, 2017.
- [20] H. Kuehne, J. Gall, and T. Serre, "An end-to-end generative framework for video segmentation and recognition," in *WACV*, 2016.
- [21] C. Lea, A. Reiter, R. Vidal, and G. D. Hager, "Segmental spatiotemporal cnns for fine-grained action segmentation," in *European Conference on Computer Vision*. Springer, 2016, pp. 36–52.
- [22] Z. Shou, J. Chan, A. Zareian, K. Miyazawa, and S.-F. Chang, "Cdc: Convolutional-de-convolutional networks for precise temporal action localization in untrimmed videos," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 5734–5743.
- [23] F. Zhou, F. De la Torre, and J. K. Hodgins, "Hierarchical aligned cluster analysis for temporal clustering of human motion," *PAMI*, 2013.
- [24] L. Rozo, S. Calinon, D. G. Caldwell, P. Jimenez, and C. Torras, "Learning physical collaborative robot behaviors from human demonstrations," *IEEE Transactions on Robotics*, vol. 32, no. 3, pp. 513–527, 2016.
- [25] N. Padoy, D. Mateus, D. Weinland, M.-O. Berger, and N. Navab, "Workflow monitoring based on 3d motion features," in *2009 IEEE 12th International Conference on Computer Vision Workshops, ICCV Workshops*. IEEE, 2009, pp. 585–592.
- [26] S. Niekum, S. Osentoski, et al., "Learning and generalization of complex tasks from unstructured demonstrations," in *IROS*, 2012.
- [27] P. Fearnhead and Z. Liu, "On-line inference for multiple changepoint problems," *Journal of the Royal Statistical Society: Series B*, 2007.
- [28] G. Konidaris, S. Kuindersma, R. Grupen, and A. Barto, "Robot learning from demonstration by constructing skill trees," *IJRR*, 2012.
- [29] B. Hayes and J. A. Shah, "Interpretable models for fast activity recognition and anomaly explanation during collaborative robotics tasks," in *ICRA*, 2017.
- [30] B. Hayes and B. Scassellati, "Autonomously constructing hierarchical task networks for planning and human-robot collaboration," in *ICRA*, 2016.