

Leaving Some Stones Unturned: Dynamic Feature Prioritization for Activity Detection in Streaming Video

Yu-Chuan Su^(✉) and Kristen Grauman

The University of Texas at Austin, Austin, USA
yosu@cs.utexas.edu

Abstract. Current approaches for activity recognition often ignore constraints on computational resources: (1) they rely on extensive feature computation to obtain rich descriptors on all frames, and (2) they assume batch-mode access to the entire test video at once. We propose a new *active* approach to activity recognition that prioritizes “what to compute when” in order to make timely predictions. The main idea is to learn a policy that dynamically schedules the sequence of features to compute on selected frames of a given test video. In contrast to traditional static feature selection, our approach continually re-prioritizes computation based on the accumulated history of observations and accounts for the transience of those observations in ongoing video. We develop variants to handle both the batch and streaming settings. On two challenging datasets, our method provides significantly better accuracy than alternative techniques for a wide range of computational budgets.

1 Introduction

Activity recognition in video is a core vision challenge. It has applications in surveillance, autonomous driving, human-robot interaction, and automatic tagging for large-scale video retrieval. In any such setting, a system that can both categorize and temporally localize activities would be of great value.

Activity recognition has attracted a steady stream of interesting research [1]. Recent methods are largely learning-based, and tackle realistic everyday activities (e.g., making tea, riding a bike). Due to the complexity of the problem, as well as the density of raw data comprising even short videos, useful video representations are often computationally intensive—whether dense trajectories, interest points, object detectors, or convolutional neural network (CNN) features run on each frame [2–8]. In fact, the expectation is that *the more features one extracts from the video, the better for accuracy*. For a practitioner wanting reliable activity recognition, then, the message is to “leave no stone unturned”, ideally extracting complementary descriptors from all video frames.

Electronic supplementary material The online version of this chapter (doi:[10.1007/978-3-319-46478-7_48](https://doi.org/10.1007/978-3-319-46478-7_48)) contains supplementary material, which is available to authorized users.

However, the “no stone unturned” strategy is problematic. Not only does it assume virtually unbounded computational resources, it also assumes that an entire video is available at once for batch processing. In reality, a recognition system will have some computational budget. Further, it may need to perform in a *streaming* manner, with access to only a short buffer of recent frames. Together, these considerations suggest some form of feature triage is needed.

Yet prioritizing features for activity in video is challenging, for two key reasons. First, the most informative features may depend critically on what has been observed so far in the specific test video, making traditional fixed/static feature selection methods inadequate. In other words, the recognition system’s belief state must evolve over time, and its priorities of which features to extract next must evolve too. Second, when processing streaming video, the entire video is never available to the algorithm at once. This puts limits on what features can even be considered each time step, and requires accounting for the feature extractors’ framerates when allocating computation.

In light of these challenges, we propose a dynamic approach to prioritize *which features to compute when* for activity recognition. We formulate the problem as policy learning in a Markov decision process. In particular, we learn a non-myopic policy that maps the accumulated feature history (state) to the subsequent feature and space-time location (action) that, once extracted, is most expected to improve recognition accuracy (reward) over a sequence of such actions. We develop two variants of our approach: one for batch processing, where we are free to “jump” around the video to get the next desired feature, and one for streaming video, where we are confined to a buffer of newly received frames. By dynamically allocating feature extraction effort, our method wisely leaves some stones *untuned*—that is, some features unextracted—in order to meet real computational budget constraints.

To our knowledge, our work is the first to actively triage feature computation for streaming activity recognition. While recent work explores ways to intelligently order feature computation in a static image for the sake of object or scene recognition [9–16] or offline batch activity detection [17], streaming video presents unique challenges, as we explain in detail below. While methods for “early” detection can fire on an action prior to its completion [18–20], they nonetheless passively extract all features in each incoming frame.

We validate our approach on two public datasets consisting of third- and first-person video from over 120 activity categories. We show its impact in both the streaming and batch settings, and we further consider scenarios where the test video is “untrimmed”. Comparisons with status quo passive feature extraction, traditional feature selection approaches, and a state-of-the-art early event detector demonstrate the clear advantages of our approach.

2 Related Work

Activity Recognition and Detection. Recognizing activities is a long-standing vision challenge [1]. Current methods explore both high-level repre-

sentations based on objects, attributes, or scenes [3, 4, 8, 21, 22], as well as holistic frame-level CNN descriptors [4–7]. Some focus on recognition in a specific domain, such as egocentric video [23, 24]. Our approach is a general algorithm for feature prioritization, and it is flexible to the descriptor type; we demonstrate instances of both types in our results. Unlike traditional activity recognition work, we account for (1) bounded computational resources for feature extraction and (2) streaming (and possibly untrimmed) input video.

Much less work addresses activity *detection*, which requires both categorizing and localizing an activity in untrimmed video, though new benchmarks aim to call more attention to the problem [25, 26]. Common strategies are sliding temporal window search [27–29] or analyzing tracked objects [30–33]. While some tracking-based methods permit incremental computation and thus can handle streaming video (e.g., [30]), they are limited to activities well-defined by a moving foreground subject. “Action-like” space-time proposals [34–37] and efficient search methods [38, 39] can avoid applying classifiers to all possible video subvolumes, but they do not prioritize feature computation. Contemporary to our work [40], a recurrent neural network that learns to predict which frame in a video to analyze next for offline action detection is proposed in [17]; its policy is free to hop forward and backward in time in the video to extract subsequent features, which is not possible in the streaming case we consider. Furthermore, our method pinpoints feature extraction requests to include not just when in the video to look for a single type of feature [17], but also where in the frame to look and which particular feature to extract upon looking there. Unlike our approach, all the above prior classifier-based methods assume batch access to the entire test video. Furthermore, with the exception of [17], they also assume features can be extracted on every frame.

Early Event Detection. The goal in “early” event detection is for the detector to fire early on in the activity instance, enabling timely reactions (e.g., for human-robot interactions [18] or nefarious activity in surveillance [19]). In [18], a structured output approach learns to recognize partial events in untrimmed video. Other methods tackle trimmed streaming video, developing novel integral-histograms that permit incremental recognition [19], or an HMM model that processes more frames until its action prediction is trusted [20]. In a sense, “early” detectors eliminate needless computation. However, the goals and methods are quite different from ours. They intend to detect an action before its completion, whereas we aim to detect an action with limited computation. As such, whereas the early methods “front-load” computation—extracting all features for each incoming frame—our method targets *which features to compute when*, and can even skip frames altogether. Furthermore, rather than learn a static model of what the onset of an action looks like, we learn a dynamic policy that indicates which computation to perform given past observations.

Fast Object Detection. Various ways to accelerate object detection have been explored [41–44]. Cascaded and coarse-to-fine detectors (e.g., [43, 44]) determine

a fixed ordering of features to quickly reject unlikely regions. In contrast, our work deals with activity recognition in video, and the feature ordering we learn is dynamic, non-myopic, and generalizes to streaming data.

Active Object and Scene Recognition in Images. Recent work considers “active” and “anytime” object recognition in images [9–16, 45, 46]. The goal is to determine which feature or classifier to apply next so as to reduce inference costs and/or supply an increasingly confident estimate as time progresses. Several methods explore dynamic feature/model selection algorithms for object and scene recognition [12, 13, 15, 16], using strategies based on reinforcement learning [11–13, 45], or myopic information gain [15, 16]. Though focused on scene recognition in images, [16] also includes a preliminary trial for “dynamic scenes” in short trimmed videos; however, the model does not represent temporal dynamics, the data is batch-processed, and gains over passive recognition are not shown. These existing methods categorize an image (recognition) [13, 46], search for an object (detection) [9–11, 14] or perform structured prediction [45].

This family of methods is most relevant to our goal. However, whereas prior work performs object/scene recognition in images, we consider activity recognition in streaming video. Feature triage on video offers unique challenges. Active recognition on images is a feature ordering task: one has the entire image in hand for processing, and the results of selected observations are static and simply accumulate. In contrast, for video, features come and go, and we must update beliefs over time and prioritize future observations accordingly. Furthermore, we must represent temporal continuity (i.e., model context over both time and space) and, when streaming, respect the hard limits of the video buffer size. In terms of a Markov decision process, this translates into a much larger state-action space.

Allocating Computation for Video. While we are not aware of any prior work that dynamically prioritizes features for streaming activity recognition, there is limited work prioritizing computation for other tasks in video. In [47], information gain is used to determine which object detectors to deploy on which frames for semantic segmentation. In [48], a second-order Markov model selects frames to apply a more expensive algorithm, for face detection and background subtraction. A cost-sensitive approach to multiscale video parsing schedules inference at different levels of a hierarchy (e.g., a group activity composed of individual actions) using AND-OR graphs [49, 50]. Aside from being different tasks than ours, all the above methods consider only the offline/batch scenario.

3 Approach

We first formalize the problem (Sect. 3.1). Then we present our approach and explain the details of its batch and streaming variants (Sect. 3.2).

3.1 Problem Formulation

Let $X \in \mathcal{X}$ denote a video clip and let $y \in \mathcal{Y}$ denote an activity category label. During training we have access to a set $\{(X_1, y_1), \dots, (X_T, y_T)\}$ of video clips, each labeled by one of L activity categories, $y_i \in \{1, \dots, L\}$. The training clips are temporally trimmed to the action of interest. At test time, we are given a novel video that may be trimmed or untrimmed. For the trimmed case, the ultimate goal is to predict the activity category label (i.e., a multi-way recognition task). For the untrimmed case, the goal is to temporally localize when an activity appears within it (i.e., a binary detection task).¹

First, we train an activity recognition module using the labeled videos. Let $\Psi(X)$ denote a descriptor computed for video X . We train an activity classifier $f: \Psi \times \mathcal{Y} \rightarrow \mathbb{R}$ to return a posterior for the specified activity category:

$$f(\Psi(X), y) = P(y|X). \quad (1)$$

We use one-vs-all multi-class logistic regression classifiers for f and bag-of-object or CNN descriptors for Ψ (details below), though other choices are possible. When training f , descriptors on training videos are fully instantiated using all frames. This classifier is trained and fixed prior to policy learning.

We formulate dynamic feature prioritization as a reinforcement learning problem: the system must learn a policy to request the features in sequence that will, over the course of a recognition *episode*, maximize its confidence, i.e. the probability estimate of f , of the true activity category. At test time, given an unlabeled video, inference is a sequential process. At each step $k = 1, \dots, K$ of an episode we must (1) actively prioritize the next feature computation *action* and (2) refine the activity category prediction. Thus, our primary goal is to learn a *dynamic policy* π that maps partially observed video features to the next most valuable action. This policy should be far-sighted, such that its choices account for interactions between the current request and subsequent features to be selected. Furthermore, it should respect a *computational budget*, meaning it conforms to constraints on the feature request costs and/or the number of inference steps permitted. We consider both *batch* and *streaming* recognition settings.

3.2 Learning the Feature Prioritization Policy

We develop a solution using a Markov decision process (MDP), which is defined by the following components [51]:

- A **state** s_k that captures the current environment at the k -th step of the episode, defined in terms of the history of extracted features and prior actions.
- A set of discrete **actions** $\mathcal{A} = \{a_m\}_{m=1}^M$ the system can perform at each step in the episode, which will lead to an update of the state. An action extracts information from the video.

¹ For clarity of presentation, in the following we present our method assuming a trimmed input video; Sect. 3.2 explains adjustments for untrimmed inputs.

- An instant **reward** $r_k = R(s_k, a^{(k)}, s_{k+1})$ received by transitioning from state s_k to state s_{k+1} after taking action $a^{(k)}$, defined in terms of activity recognition. The total reward is $\sum_k \gamma^k R(s_k, a^{(k)}, s_{k+1})$, where $\gamma \in [0, 1]$ is a discount factor on future rewards. Larger values lead to more far-sighted policies.
- A **policy** $\pi : s \rightarrow a$ determines the next action based on the current state. It selects the action that maximizes the expected reward:

$$\pi(s_k) = \underset{a}{\operatorname{argmax}} E[R|s_k, a, \pi], \quad (2)$$

for this action and future actions continuing under the same policy.

We next detail the video representation, state-action features, and rewards for the general case. Then, we define aspects specific to the batch and streaming settings, respectively.

Video Descriptors and Actions. Our algorithm accommodates a range of descriptor/classifier choices. The requirements are that the descriptor (1) have temporal locality, and (2) permit incremental updates as new descriptor instances are observed. These specs are met by popular “bag-of- X ” and CNN frame features, as we will demonstrate in results, as well as others like quantized dense trajectories or human body poses.

We focus our implementation primarily on a *bag-of-objects* descriptor. Suppose we have object detectors for N object categories. The fully observed descriptor $\Psi(X)$ is an N -dimensional vector, where $\Psi_n(X)$ is the likelihood that the n -th object appears (at least once) in the video clip X . We chose a bag-of-objects for its strength in compactly summarizing high-level content relevant to activities [3, 4, 52]. For example, an activity like “making sandwich” is definable by bread, knife, frig, etc. Furthermore, it exposes semantic temporal context valuable for sequential feature selection. For example, after seeing a mug, the system may learn to look next for either a tea bag or a coffee maker.

Each step in an episode performs some action $a^{(k)} \in \mathcal{A}$ at a designated time t^k in the video. We define each action as a tuple $a_m = \langle o_m, l_m \rangle$ consisting of an object and video location.² Specifically, $o_m \in \{1, \dots, N\}$ specifies an object detector, and l_m specifies the space-time subvolume where to run it. The observation result x_m of taking action a_m is the maximum detection probability of object o_m in volume l_m .³ It is used to incrementally refine the video representation $\Psi(X)$. Let $o^{(k)} = n$ denote the object specified by selected action $a^{(k)}$. Upon receiving $x^{(k)}$, the n -th entry in $\Psi(X^k) \in \mathbb{R}^N$ is updated by taking the maximum observed probability for that object so far:

$$\Psi_n(X^k) = \max \left(\Psi_n(X^{k-1}), x^{(k)} \right), \quad (3)$$

² Note that $a^{(k)}$ identifies an action selected at step k in the episode, whereas a_m is one of the M discrete action choices in \mathcal{A} .

³ Some object detectors share features across object categories, e.g., R-CNN [53], in which case it may be practical to simplify the action to select only the video volume and apply all object classes. We use the DPM detector [54], which has the advantage of near real-time detection [42] using a single thread, whereas R-CNN relies heavily on parallel computation and hardware acceleration [55].

where $\Psi(X^k)$ denotes the video representation based on the observation results up to the k -th step of the episode. The initialization of $\Psi(X)$ is explained below.

To alternatively apply our method with CNN features—which show promise for video (e.g., [5–7])—we define the representation and actions as follows. The video representation averages per-frame CNN descriptors:

$$\Psi_n(X^k) = \text{mean}(X^k), \quad (4)$$

and the action becomes $a_m = l_m$, since we need to specify the temporal location alone. Though very fast CNN extraction is possible (76 fps on a CPU [56]), conventional approaches still require time linear in the length of the video, since they touch each frame. We offer *sub-linear* time extraction; for example, our results maintain accuracy for streaming recognition with CNNs while pulling the features from fewer than 1% of the frames.

State-Action Features. With Q-learning [51], the value of actions $E[R|s, a, \pi]$ in Eq. (2) is evaluated with $Q^\pi(s, a)$. It must return a value for any possible state-action pair. Our state space is very large—equal to the number of possible features times the number of possible space-time locations times their possible output values. This makes exact computation of $Q^\pi(s, a)$ infeasible. Thus, as common in such complex scenarios, we adopt a linear function approximation $Q^\pi(s, a) = \theta^T \phi(s, a)$, where $\phi(s, a)$ is a feature representation of a *state-action* pair and θ is learned from activity-labeled training clips (explained below).

The state-action feature $\phi(s, a)$ encodes information relevant to policy learning: the *previous object detection* results and the *action history*. Past object detections help the policy learn to exploit object co-occurrences (e.g., that running a laptop detector after finding soap is likely wasteful) and select discriminative but yet-unseen objects (e.g., having seen a chair, looking next for a bed or dish could disambiguate the bedroom or kitchen context, whereas a cell phone would not). The action history can also benefit the policy, letting it learn to avoid redundant selections.

Motivated by these requirements, we define the state-action feature $\phi(s, a) \in \mathbb{R}^{N+M}$ as

$$\phi(s_k, a) = [\Psi(X^k), \delta t^k], \quad (5)$$

where $\Psi(X^k)$ encodes the detection results and δt^k encodes the action history. $\Psi(X^k) \in \mathbb{R}^N$ is the representation defined above. The action history feature $\delta t^k \in \mathbb{R}^M$ encodes how long it has been since each action was performed in the episode, which for action m is

$$\delta t^k(m) = t^k - \max_i \{t^i | a^{(i)} = a_m\}, \quad (6)$$

with $\delta t^k(m) = 0$ if a_m has never been performed before.

To encode actions into the state-action representation $\phi(s, a)$, we learn one linear model θ_{a_m} for each action (details below), such that $Q^\pi(s, a_m) = \theta_{a_m}^T \phi(s, a)$. In the following, we denote $\theta = \{\theta_{a_m}\}_{m=1}^M$.

Reward. We define a smooth reward function that rewards increasing confidence in the correct activity label, our ultimate prediction task. Intuitively, the model

should continuously gather evidence for the activity during the episode, and its confidence in the correct label should increase over time and surpass all other activities by the time the computation budget is exhausted. Accordingly, for a training episode run on video X with label y^* , we define the reward:

$$R(s_k, a^{(k)}, s_{k+1}) = f(\Psi(X^{k+1}), y^*) - f(\Psi(X^k), y^*). \tag{7}$$

With this definition, a new action gets no “credit” for confidence attributable to previous actions. We found that rewarding accuracy increases per unit time performs similarly to training multiple policies targeting fixed budgets. Moreover, the proposed reward has the advantage that we can run the policy for as long as desired at test time, which is essential for streaming video. Fixed-budget policies, though common in RL, are ill-suited for streaming data since we cannot know in advance the test video’s duration and the budget to allocate.

Dynamic Feature Prioritization Policy. We learn the policy π using policy iteration [51]. Policy iteration is an iterative algorithm that alternates between learning θ from samples $\{Q^\pi(s, a), \phi(s, a)\}$ and generating samples using the learned π . Given the policy parameter $\theta^{(i)}$ at iteration i , we generate samples by running recognition episodes on training videos and collect the state features $\phi(s_k, a^{(k)})$ and instant rewards r_k from each step of the episode. The action reward $Q^\pi(s_k, a^{(k)})$ is the total reward from Eq. 7 (with discounting) after finishing the episode. See Supp. for details.

After collecting the training samples, the new policy parameter $\theta^{(i+1)}$ is learned by solving one ridge regression for each action. The algorithm then iterates, generating new samples using $\theta^{(i+1)}$. We run a fixed number of iterations to learn the policy and keep all the training samples from previous iterations when learning the new policy parameters. To ensure the algorithm sufficiently explores the state space, we use an ϵ -greedy algorithm, selecting a random action instead of $a^{(k)} = \pi(s_k)$ with probability ϵ when generating training samples.

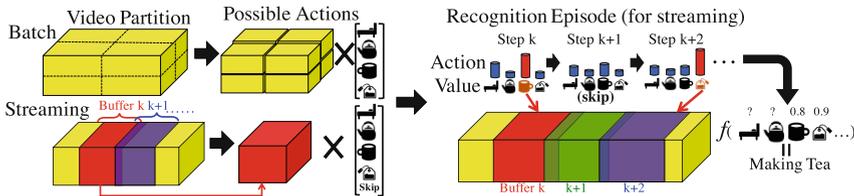


Fig. 1. Action spaces. **Left top:** In batch, the whole video is divided into subvolumes, and actions are defined by the volume and object category to detect. **Left bottom:** In streaming, the video is divided into segments by the buffer at each step, and actions are the object category to detect in the buffer plus a “skip” action. **Right:** Our method learns a policy to dynamically select a sequence of useful features to extract. Note this episode depicts the streaming case.

Batch Recognition Setting. In the *batch recognition* setting, we have access to the entire test video throughout an episode, and the budget is the total resources available for feature computation, i.e., as capped by episode length K . In this case, our model is free to run an object detector at arbitrary locations. Most existing activity recognition work assumes this setting, though without imposing a computation limit. It captures the situation where one has an archive of videos to be recognized offline, subject to real-world resource constraints (e.g., auto-tagging YouTube clips under a budget of CPU time).

Each candidate location l_m in the action set is a spatio-temporal volume. Its position and size is specified relative to the entire clip, so that the number of possible actions is constant even though video lengths and resolutions may vary. We use non-overlapping volumes splitting the video in half in each dimension. See Fig. 1, top left. Note that while the bag-of-objects discards order, the action set *preserves* it. That means our policy can learn to exploit the space-time layout of objects if/when beneficial to feature prioritization (e.g., learning it is useful to look for a washing machine *after* a laundry basket, or an pot *above* a stove).

In the batch setting, performing the same action at different steps in the episode will produce the same observation. Without loss of generality, we define the time an action is performed as a constant $t^k = \text{const}.\forall k$, and the action history feature δt^k becomes a binary indicator showing whether an action has been performed in the episode. We forbid the policy to choose actions that have been performed since they provide no new information.

By design, the bag-of-objects is accumulated over time. We impute the observations of un-performed actions by exploiting previously learned object co-occurrence statistics. In particular, we represent the M -dimensional distribution over the action space with a Gaussian Mixture Model (GMM). We learn its parameters on the same data that trains f , which has full object detection results on all videos. We initialize $\Psi(X^0)$ with the average posteriors observed in that same training set. Then to impute \tilde{x}_u for an un-taken action, we take the expected value using the GMM.

Streaming Recognition Setting. In the *streaming* setting, recognition takes place at the same time the video stream is received, so the model can only access frames received before the current time step. Further, the model has a fixed size buffer that operates in a first-in-first-out manner; its feature requests may only refer to frames in the current buffer. Though largely unexplored for activity recognition, the streaming scenario is critical for applications with stringent resource constraints. For example, when capturing long-term surveillance video or wearable camera data, it may be necessary to make decisions online without storing all the data.

The feature extractor can process a fixed number of frames per second (FPS), and this rate indirectly determines the resource budget. That is, the faster the feature extractors can run, the more of them we can apply as the buffer moves forward. A recognition episode ends when it reaches the end of a video stream.

The action space consists of the N object detectors (or alternatively, the single CNN descriptor); an action’s space-time location l_m is always the entire current buffer. We further define a *skip* action a_0 , which instructs the model to wait until the next frame arrives without performing any feature extraction. Thus, for streaming, the number of actions equals the number of objects plus one ($M = N + 1$). See Fig. 1, bottom left. The skip action saves computation when the model expects a new observation will not benefit the recognition task. For example, if the model is confident that the video is taken in a bedroom, and all un-observed objects would appear only in the bathroom, then forcing the system to detect new objects is wasteful.

Because new frames may arrive and old frames may be discarded during an action, the video content available to the model will change between steps; performing the same action at different steps yields different observations. To connect the video content in the buffer and the actions in the episode, we define the time t^k of the k -th action using the last frame number in the buffer when the action was issued by the policy.

While we assume so far the video contains only the target activity, i.e. the video is trimmed to the span of the activity, our method generalizes to *untrimmed activity detection* in the streaming environment. In that case, the target activity only occurs in part of the video, and the system must identify the span where the activity happens. This is non-trivial in the streaming environment.

To handle the streaming input, we pose the problem in terms of frame-level labeling: we predict a label for each frame as it is received, and the activity detector must optimize accuracy across all frames. However, we do not estimate the activity label from a single frame alone. Rather, we predict each frame’s label using the temporal window around it. For every newly arrived frame, we consider all the windows shorter than an upper bound β that end at the frame. We predict the label of each window based on the same representation as trimmed video, and we select the one with highest confidence as the prediction result of the target frame. Note that this requires storing only the descriptors for recent history of length β , but keeping no video beyond the current buffer. The activity recognizer f is a binary classifier trained to determine whether the target activity occurs in the window, and actions are terminated when a new frame arrives. Whereas non-streaming methods can utilize complete sliding windows directly to predict the activity span offline [25,26], our method aggregates its shorter streaming temporal window results to produce full detection windows.

4 Experiment

Datasets. We evaluate on two datasets: the Activities of Daily Living [3] (ADL) and UCF-101 [57]. **ADL** consists of 313 egocentric videos recorded by 14 subjects, labeled with $L = 18$ activity categories (e.g., making coffee, using computer). Following [3], we train f in a leave-one-subject-out manner. Our policy is learned on a disjoint set of 110 clips (those used in [3] for training object detectors). As observations $x^{(k)}$, we use the provided object detector outputs for

$N = 26$ categories (1 fps). **UCF-101** consists of 13,320 YouTube videos covering $L = 101$ activities. We use the provided training splits to train f , reserve half of the test splits for policy learning, and average results over all 6 splits. As observations $x^{(k)}$, we use the object detector outputs for $N = 75$ objects, kindly shared by the authors of [4], which are frame-level scores (no bbox).⁴ For CNN frame descriptors, we use the fc-7 activation of VGG-16 [58] from Caffe Model Zoo (1 fps). The video clips average 78 and 19s for ADL and UCF, respectively.

To create test data and policy learning data for the *untrimmed* experiments, we concatenate multiple clips following [18,38]. See Supp. for details. In all, we obtain 8,410 (UCF) and 3,130 (ADL) untrimmed sequences, with lengths averaging 2–7 min.

Baselines. We compare to several methods:

- **Passive:** selects the next action randomly. It represents the most direct mapping of existing activity recognition methods to the resource-constrained regime. The system does not actively decide which features to extract.
- **Object-Preference [4]:** a static feature selection heuristic employed for bag-of-objects activity recognition. It prioritizes objects that appear frequently in each activity. We average x_m per activity and order a_m based on its maximum response over all activities. Though the authors intend this metric to identify the most discriminative objects—not to sequence feature extraction—it is a useful reference point for how far one can get with static feature selection.
- **Decision tree (DT):** a static feature ordering method. We learn a DT to recognize activities, where the attribute space consists of the Cartesian product of object detectors and subvolume locations (l_m). We sort the selected attributes by their Gini importance [59]. In the streaming case, we test two variations: DT-Static, where we cycle through the features in that order, and DT-Top, where we take only the top P features and repeatedly apply all those object detectors on each frame. P is equal to the object detector framerate. Thus, DT-Top runs as many detectors as it can at framerate, prioritizing those expected to be most discriminative.
- **Max-Margin Early Event Detector (MMED) [18]:** a state-of-the-art early event detector designed for untrimmed streaming video. It aims to fire on the activity as soon as possible after its onset. We implement it based on structure SVM solver BCFW [60] and apply the authors’ default parameter settings. The same window search process as in the untrimmed variant of our method is used for prediction, with a window size ranging from 1 to β frames.

Implementation Details. Please see supplementary material.

Streaming Activity Recognition. First we test the streaming setting. In this case, feature extraction speed (e.g., object detector speed) dictates the action budget: the faster the features can be extracted, the more can be used while

⁴ We retain the 75 objects among all 15,000 found most responsive for the activities, following [4]. Because the provided detections are frame-level, we split volumes only in the temporal dimension for l_m on UCF.

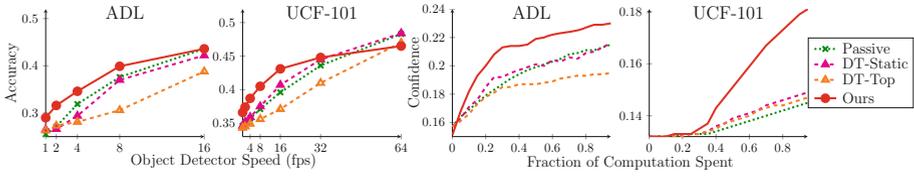


Fig. 2. Streaming recognition result. Left: Recognition accuracy as a function of object detector speed. Our method reduces computation by more than 50 % under the same accuracy. Right: Confidence score improvement as the episode progresses. Our method improves the prediction more rapidly than the baselines, indicating that it selects more informative features.

keeping up with the incoming video framerate. We stress that to our knowledge, no prior activity recognition work considers feature triage for streaming video.

Figure 2 (left 2 plots) shows the final recognition accuracy at the episode’s completion, as a function of the object detectors’ speed.⁵ Our method performs better than the rest, across the range of detector speeds. Overall, our method reduces cost by 80 % and 50 % on UCF and ADL, respectively. The left side of the plots is most interesting; by definition all methods will converge in accuracy once the object detector framerate equals the number of possible objects to detect (26 for ADL and 75 for UCF). DT-Top is the weakest method for this task. It repeatedly uses only the most informative features, but they are insufficient to discriminate the 18 to 101 different activities. This result shows the necessity of instance-dependent feature selection, which our method provides. Because our method can skip frames, the actual amount of computation spent does not grow linearly with detector speed. So, though hidden in Fig. 2, our method performs much less computation at the higher detector speeds. Figures 3 and 4 will more directly show our runtime advantage.

Figure 2 (right 2 plots) shows the confidence score (of the ground truth activity) over the course of the episodes. Here we apply the 8 fps detector. The

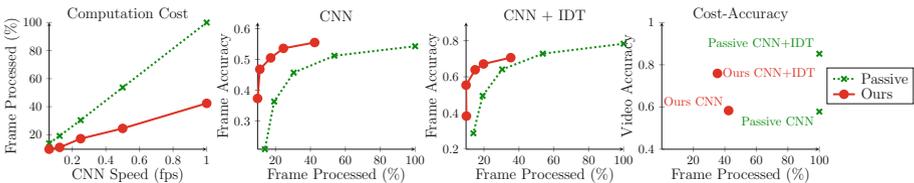


Fig. 3. Streaming recognition result on UCF-101 using CNN frame features. Our method achieves over 90 % of the ultimate accuracy by processing only 20 % of sampled frames. With the 1 fps sample rate, this corresponds to 0.8 % of frames in the entire video.

⁵ Object-Pref [4] is not applicable to the streaming case because it lacks a unique object response prior for the actions that is dependent on the buffer location.

baseline methods improve their prediction smoothly, which indicates that they collect meaningful detection results at the same rate throughout the episode. In contrast, our method begins to improve rapidly after some point in the episode. This shows that it starts to collect more useful information once it has explored the novel video sufficiently. Because UCF uses about $4\times$ more objects in the representation, it takes more computation (actions) before the representation converges. See Supp. for qualitative analysis of the policies learned.

Figure 3 shows our method has clear advantages if applied with CNN features as well.⁶ Here the DT baselines are not applicable, since there is only one feature type; the question is whether to extract it or not. The Passive baseline uniformly distributes its frame selections. The left plot shows that no matter the framerate of the CNN extractor, our method requires less than half of the frames to achieve the same accuracy. The second plot shows our method achieves peak accuracy looking at just a fraction of the streaming frames, where the accuracy is measured over every step in the recognition. Our algorithm skips 80% of the frames, but still achieves over 90% of the ultimate accuracy obtained using *all* frames. With the base sampling rate of 1 fps, processing 20% of the frames means we extract features for only 0.8% of the entire video.

In the third plot, we further combine improved dense trajectories [2] (IDT) with the CNN features to show that our method can benefit from more powerful features without modification. The right plot compares the cost-accuracy tradeoff between the ultimate multi-class accuracy achieved by our streaming method vs. that attained using exhaustive feature extraction. We obtain similar accuracies with substantially less computation.

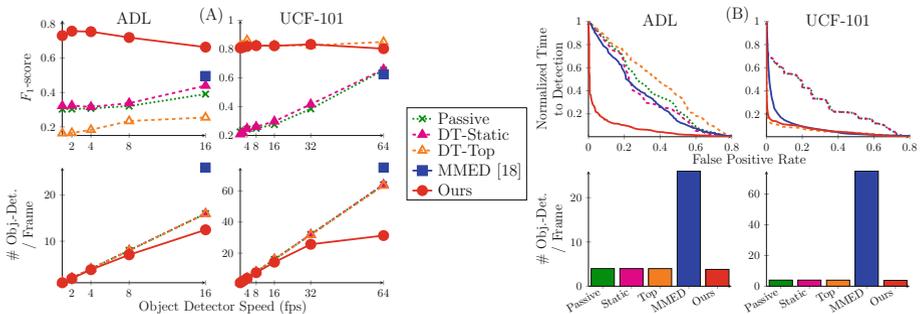


Fig. 4. Streaming untrimmed detection results, with comparison to [18]. **(A):** Accuracy (top, higher is better) and computation cost (bottom) as a function of object detector framerate. **(B):** Activity monitoring operating curves (top, lower is better) and corresponding computational costs (bottom).

⁶ ADL is less amenable to full-frame CNN descriptors, due to domain shift of egocentric video and the nature of the composite, object-driven activities.

Untrimmed Video Activity Detection. Next we evaluate streaming detection for untrimmed video. This setting permits comparison with the state of the art MMED [18] “early” activity detector.

Since we must predict whether each frame is encompassed by the target activity, we measure accuracy with the F_1 -score. While we assume the episode terminates after reaching the end of the video stream in our algorithm, in some applications it may be sufficient to identify the occurrence of the activity and then terminate the episode. Therefore, we further compare the detection timeliness using the Activity Monitoring Operating Curve (AMOC), following [18]. AMOC is the normalized time to detection (NT2D) vs. the false positive rate curve. The lower the value, the better the timeliness of the detector.

In Fig. 4(A), the top plots show the F_1 -scores. Overall, our method performs the best in terms of accuracy. On ADL, we achieve nearly twice the accuracy of all baselines until the object detector speed reaches 16 fps. On UCF, our method is comparable to the best baseline, DT-Top. Whereas DT-Top is weak on UCF for the multi-class recognition scenario (see above), it fares well for binary detection on this dataset. This is likely because the UCF activities are often discriminated by one or few key objects, and we give the baselines the advantage of pruning the object set to those most responsive on each activity.

The bottom two plots in Fig. 4(A) show the actual number of object detectors run. Our method reduces computation cost significantly under high object detector speeds, thanks to its ability to forgo computation with the “skip” action. In particular, it performs 50% fewer detections under 64 fps on UCF while maintaining accuracy. On the other hand, the baseline methods’ cost grows linearly with the object detector speed.

Figure 4(B) shows the AMOC under 4 fps detection speed (top, see Supp. for others) and the associated computational costs (bottom). Despite the fact our reward function does not specifically target this metric, our method achieves excellent timeliness in detection. MMED performs second best on the metric, but it incurs much higher computation cost than ours, as shown by the bar charts. This is because MMED is trained to fire early, but always extracts all features in the frames it does process.

Batch Activity Recognition. Finally, we test the batch setting. We evaluate accuracy as a function of the computation budget—the fraction of all possible actions the algorithm performs (i.e., the number of features it extracts, normalized by video length). “All possible” features would be extracting all features in all frames (1 fps).

Figure 5 shows the results. Our method outperforms the baselines, especially when the computation budget is low (<0.5). In fact, extracting only 30% of the features on ADL, we achieve the same accuracy as with all features. Without a budget constraint, the video representation will converge to that of the full observation—no matter what method is used; that is, all methods must attain the same accuracy on the rightmost point on each plot. Our method shows more significant gains on ADL than UCF. We think this reflects the fact that the object categories for ADL are tailored well for the activities (e.g., household

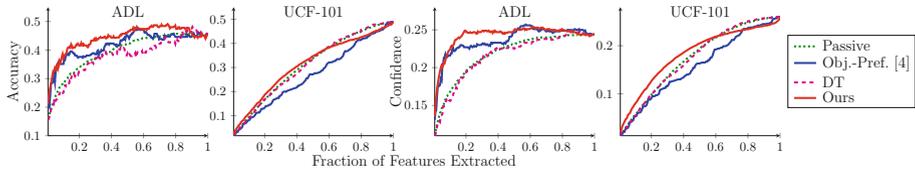


Fig. 5. Batch recognition accuracy/confidence score vs. computation budget. Our method outperforms the baselines, especially when the computation budget is low.

items), whereas the object bank for UCF includes arbitrary objects which may not even appear in the dataset. Furthermore, ADL has more objects in any single activity, offering more signal for our method to learn. Object-Pref [4] is next best on ADL, though it is noticeably weaker on UCF because it does not account for the temporal redundancy of the dataset. Our method is 2.5 times faster than this nearest competing baseline.

Surprisingly, the Decision Tree (DT) baseline performs similarly to Passive. (Note that DT-Static only is used; DT-Top is applicable only for the streaming case.) We attempted to improve its accuracy by learning it on the same features as f , i.e., dropping the subvolumes from the attributes and running one object detector over the entire video for each action. However, this turned out to be worse due to redundant/wasteful detections. This shows the importance of coping with partially observed results, which the proposed method can do.

Our contribution is not a new model for activity recognition, but instead a method that enables activity recognition for existing features/classifiers without exhaustive feature computation. This means the accuracies achieved with “all features” is the key yardstick to hold our results against. Nonetheless, to put in context with other systems: the base batch recognition model we employ gets results slightly better than the state-of-the-art on ADL [3, 61] and within 4.5–11 state-of-the-art using comparable features on UCF [4, 5].

5 Conclusions

We developed a dynamic feature extraction strategy for activity recognition under computational constraints. On two diverse datasets, our method shows competitive recognition performance under various resource limitations. It can be used to consistently achieve *better accuracy* under the same resource constraint, or meet a given accuracy using *less resources*. In future work we plan to investigate policies that reason about variable cost descriptors.

Acknowledgement. This research is supported in part by ONR PECASE N00014-15-1-2291. We thank Dr. Mihir Jain for sharing the object detection results on UCF101 and the detailed information of the data.

References

1. Aggarwal, J.K., Ryoo, M.S.: Human activity analysis: a review. *ACM Comput. Surv.* **43**(3), April 2011
2. Wang, H., Schmid, C.: Action recognition with improved trajectories. In: *ICCV* (2013)
3. Pirsiavash, H., Ramanan, D.: Detecting activities of daily living in first-person camera views. In: *CVPR* (2012)
4. Jain, M., van Gemert, J.C., Snoek, C.G.M.: What do 15,000 object categories tell us about classifying and localizing actions? In: *CVPR* (2015)
5. Simonyan, K., Zisserman, A.: Two-stream convolutional networks for action recognition in videos. In: *NIPS* (2014)
6. Xu, Z., Yang, Y., Hauptman, A.: A discriminative cnn video representation for event detection. In: *CVPR* (2015)
7. Zha, S., Luisier, F., Andrews, W., Srivastava, N., Salakhutdinov, R.: Exploiting image-trained CNN architectures for unconstrained video classification. In: *BMVC* (2015)
8. Han, D., Bo, L., Sminchisescu, C.: Selection and context for action recognition. In: *ICCV* (2009)
9. Butko, N., Movellan, J.: Optimal scanning for faster object detection. In: *CVPR* (2009)
10. Vijayanarasimhan, S., Kapoor, A.: Visual recognition and detection under bounded computational resources. In: *CVPR* (2010)
11. Karayev, S., Baumgartner, T., Fritz, M., Darrell, T.: Timely object recognition. In: *NIPS* (2012)
12. Dulac-Arnold, G., Denoyer, L., Thome, N., Cord, M.: Sequentially generated instance-dependent image representations for classification. In: *ICLR* (2014)
13. Karayev, S., Fritz, M., Darrell, T.: Anytime recognition of objects and scenes. In: *CVPR* (2014)
14. Gonzalez-Garcia, A., Vezhnevets, A., Ferrari, V.: An active search strategy for efficient object class detection. In: *CVPR* (2015)
15. Gao, T., Koller, D.: Active classification based on value of classifier. In: *NIPS* (2011)
16. Yu, X., Fermuller, C., Teo, C.L., Yang, Y., Aloimonos, Y.: Active scene recognition with vision and language. In: *CVPR* (2011)
17. Yeung, S., Russakovsky, O., Mori, G., Fei-Fei, L.: End-to-end learning of action detection from frame glimpses in videos. In: *CVPR* (2016)
18. Hoai, M., la Torre, F.D.: Max-margin early event detectors. In: *CVPR* (2012)
19. Ryoo, M.: Human activity prediction: early recognition of ongoing activities from streaming videos. In: *ICCV* (2011)
20. Davis, J., Tyagi, A.: Minimal-latency human action recognition using reliable-inference. *Image Vis. Comput.* **24**, 455–472 (2006)
21. Yao, B., Jiang, X., Khosla, A., Lin, A., Guibas, L., Fei-Fei, L.: Human action recognition by learning bases of action attributes and parts. In: *ICCV* (2011)
22. Rohrbach, M., Regneri, M., Andriluka, M., Amin, S., Pinkal, M., Schiele, B.: Script data for attribute-based recognition of composite activities. In: Fitzgibbon, A., Lazebnik, S., Perona, P., Sato, Y., Schmid, C. (eds.) *ECCV 2012*. LNCS, vol. 7572, pp. 144–157. Springer, Heidelberg (2012). doi:[10.1007/978-3-642-33718-5_11](https://doi.org/10.1007/978-3-642-33718-5_11)
23. Li, Y., Ye, Z., Rehg, J.M.: Delving into egocentric actions. In: *CVPR* (2015)

24. Ma, M., Fan, H., Kitani, K.M.: Going deeper into first-person activity recognition (2016)
25. Jiang, Y.G., Liu, J., Roshan Zamir, A., Toderici, G., Laptev, I., Shah, M., Sukthankar, R.: THUMOS challenge: action recognition with a large number of classes (2014). <http://csrcv.ucf.edu/THUMOS14/>
26. Heilbron, F.C., Escorcia, V., Ghanem, B., Nibbles, J.C.: ActivityNet: a large-scale video benchmark for human activity understanding. In: CVPR (2015)
27. Ke, Y., Sukthankar, R., Hebert, M.: Efficient visual event detection using volumetric features. In: ICCV (2005)
28. Duchenne, O., Laptev, I., Sivic, J., Bach, F., Ponce, J.: Automatic annotation of human actions in video. In: ICCV (2009)
29. Sathkin, S., Hebert, M.: Modeling the temporal extent of actions. In: Daniilidis, K., Maragos, P., Paragios, N. (eds.) ECCV 2010. LNCS, vol. 6311, pp. 536–548. Springer, Heidelberg (2010). doi:[10.1007/978-3-642-15549-9_39](https://doi.org/10.1007/978-3-642-15549-9_39)
30. Medioni, G., Nevatia, R., Cohen, I.: Event detection and analysis from video streams. *Trans. Pattern Anal. Mach. Intell.* **23**(8), 873–889 (2001)
31. Yao, A., Gall, J., van Gool, L.: A hough transform-based voting framework for action recognition. In: CVPR (2010)
32. Kläser, A., Marszałek, M., Schmid, C., Zisserman, A.: Human focused action localization in video. In: International Workshop on Sign, Gesture, Activity (2010)
33. Lan, T., Wang, Y., Mori, G.: Discriminative figure-centric models for joint action localization and recognition. In: ICCV (2011)
34. Yu, G., Yuan, J.: Fast action proposals for human action detection and search. In: CVPR (2015)
35. Jain, M., van Gemert, J., Jegou, H., Bouthemy, P., Snoek, C.: Action localization with tubelets from motion. In: CVPR (2015)
36. Gkioxari, G., Malik, J.: Finding action tubes. In: CVPR (2015)
37. van Gemert, J.C., Jain, M., Gati, E., Snoek, C.G.M.: APT: action localization proposals from dense trajectories. In: BMVC (2015)
38. Chen, C.Y., Grauman, K.: Efficient activity detection with max-subgraph search. In: CVPR (2012)
39. Yu, G., Yuan, J., Liu, Z.: Unsupervised random forest indexing for fast action search. In: CVPR (2011)
40. Su, Y.C., Grauman, K.: Leaving some stones unturned: dynamic feature prioritization for activity detection in streaming video. Technical report, Comp. Sci. Dept., Univ. Texas, Austin, AI-15-05 (2015)
41. Sadeghi, M.A., Forsyth, D.: 30Hz object detection with DPM V5. In: Fleet, D., Pajdla, T., Schiele, B., Tuytelaars, T. (eds.) ECCV 2014. LNCS, vol. 8689, pp. 65–79. Springer, Heidelberg (2014). doi:[10.1007/978-3-319-10590-1_5](https://doi.org/10.1007/978-3-319-10590-1_5)
42. Yan, J., Lei, Z., Wen, L., Li, S.: The fastest deformable part model for object detection. In: CVPR (2014)
43. Viola, P., Jones, M.: Rapid object detection using a boosted cascade of simple features. In: CVPR (2001)
44. Pedersoli, M., Vedaldi, A., Gonzalez, J.: A coarse-to-fine approach for fast deformable object detection. In: CVPR (2011)
45. Weiss, D.J., Taskar, B.: Learning adaptive value of information for structured prediction. In: NIPS (2013)
46. Wang, J., Bolukbasi, T., Trapeznikov, K., Saligrama, V.: Model selection by linear programming. In: Fleet, D., Pajdla, T., Schiele, B., Tuytelaars, T. (eds.) ECCV 2014. LNCS, vol. 8690, pp. 647–662. Springer, Heidelberg (2014). doi:[10.1007/978-3-319-10605-2_42](https://doi.org/10.1007/978-3-319-10605-2_42)

47. Karasev, V., Ravichandran, A., Soatto, S.: Active frame, location, and detector selection for automated and manual video annotation. In: CVPR (2014)
48. Chen, D., Bilgic, M., Getoor, L., Jacobs, D.: Dynamic processing allocation in video. *IEEE Trans. Pattern Anal. Mach. Intell.* **33**(11), 2174–2187 (2011)
49. Amer, M.R., Xie, D., Zhao, M., Todorovic, S., Zhu, S.-C.: Cost-sensitive top-down/bottom-up inference for multiscale activity recognition. In: Fitzgibbon, A., Lazebnik, S., Perona, P., Sato, Y., Schmid, C. (eds.) *ECCV 2012. LNCS*, vol. 7575, pp. 187–200. Springer, Heidelberg (2012). doi:[10.1007/978-3-642-33765-9_14](https://doi.org/10.1007/978-3-642-33765-9_14)
50. Amer, M., Todorovic, S., Fern, A., Zhu, S.C.: Monte Carlo tree search for scheduling activity recognition. In: *ICCV* (2013)
51. Russell, S., Norvig, P.: *Artificial Intelligence: A Modern Approach*. Pearson, Upper Saddle River (2010)
52. Gupta, A., Davis, L.S.: Objects in action: an approach for combining action understanding and object perception. In: *CVPR* (2007)
53. Girshick, R., Donahue, J., Darrell, T., Malik, J.: Rich feature hierarchies for accurate object detection and semantic segmentation. In: *CVPR* (2014)
54. Felzenszwalb, P.F., Girshick, R.B., McAllester, D., Ramanan, D.: Object detection with discriminatively trained part based models. *PAMI* **32**(9), 1627–1645 (2010)
55. Ren, S., He, K., Girshick, R., Sun, J.: Faster R-CNN: towards real-time object detection with region proposal networks. In: *NIPS* (2015)
56. Nvidia: GPU-based deep learning inference: a performance and power analysis. *Whitepaper* (2015)
57. Soomro, K., Zamir, A.R., Shah, M.: UCF101: A dataset of 101 human actions classes from videos in the wild. *arXiv preprint [arXiv:1212.0402](https://arxiv.org/abs/1212.0402)* (2012)
58. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. *arXiv preprint [arXiv:1409.1556](https://arxiv.org/abs/1409.1556)* (2014)
59. Breiman, L.: Random forests. *Mach. Learn.* **45**(1), 5–32 (2001)
60. Lacoste-Julien, S., Jaggi, M., Schmidt, M., Pletscher, P.: Block-coordinate Frank-Wolfe optimization for structural SVMs. In: *ICML* (2013)
61. McCandless, T., Grauman, K.: Object-centric spatio-temporal pyramids for ego-centric activity recognition. In: *BMVC* (2013)