

# Duration Dependent Codebooks for Change Detection

Brandon A. Mayer  
 Brandon\_Mayer@brown.edu  
 Joseph L. Mundy  
 mundy@lems.brown.edu

Brown University  
 School of Engineering  
 Rhode Island, USA

Change detection is a computer vision application that attempts to distinguish normal and abnormal scene activity in video sequences. However, natural scenes are composed of complex, dynamic events that make it difficult for a change detection system to distinguish between changes of interest and background. To further compound the problem, it is impossible to define what a system should consider as a relevant change without considering the context of the application. For example, are cars moving along a highway foreground or background? If the goal of the application is to count the number of cars entering and exiting a restricted area, it is necessary for the system to account for every car in the scene. However, if the requirement is to monitor a busy highway for irregular traffic activity such as a collision, then the system will need to consider common traffic patterns as normal and not declare routine traffic activity as change.

This paper describes a supervised system for pixel-level change detection for fixed, monocular surveillance cameras. Per-pixel intensity sequences are modeled by a class of Hidden Semi-Markov Models, Duration Dependent Hidden Markov Models (DDHMMs), to accurately account for stochastically periodic phenomena prevalent in real-world video. The per-pixel DDHMMs are used to assign discrete state labels to pixel intensity sequences that summarize the appearance and temporal statistics of the observations. State assignments are then used as a features for constructing per-pixel code books during a training phase to identify changes of interest in new video.

The per-pixel intensity model is validated by showing superior predictive performance to pixel representations commonly used in change detection applications. A new data set is presented which contain dynamic, periodic backgrounds with larger time scale variability than previous data sets and the proposed method is compared to state-of-the-art change detection methods using the new videos.




Figure 1: Graphical visualization of the DDHMM

A Duration Dependent Hidden Markov Model (DDHMM) models a sequence of observations,  $Y = (y_1, y_2, \dots, y_T)$ , using a sequence of latent state pairs:  $((S_1, D_1), (S_2, D_2), \dots, (S_l, D_l))$  where  $S_i$  is a state label and  $D_i$  is a random variable that represents the time spent in state  $S_i$ . Note that capital letters denote random variables and lower case letters represent specific variable assignments. A graphical visualization of a DDHMM is shown in Figure 1 where dotted circles represent random variables and the shaded nodes represent observed quantities. The topology of the graphical model is variable since the number of state-duration tuples will change depending on the particular configuration of the duration random variables.

The observation and state sequences are related through three fundamental distributions: the duration  $p(D_i = d_i | S_i = s_i)$ , state transition  $p(S_i = s_i | S_{i-1} = s_{i-1})$  and emission  $p(y_t | S_i = s_i)$  distributions. The likelihood of an observation sequence given a particular latent state assignment is given by equation 1 where  $r_i = \sum_{m=1}^i d_m$  and  $p(s_1)$  is an initial distribution of state labels. The observation sequence is assumed to be left-censored, i.e., the last tuple  $(s_l, d_l)$  is distributed according to the state survival distribution  $p(D_l \geq d_l | s_l)$ , to mitigate the effect of the length of the observation sequence on the probability of a particular state

sequence [2].

$$p(y_1, \dots, y_T | (s_1, d_1), \dots, (s_l, d_l)) = \dots$$

$$p(s_1) p(d_1 | s_1) \prod_{m=1}^{d_1} p(y_m | s_1) \prod_{i=2}^{l-1} p(d_i | s_i) p(s_i | s_{i-1}) \dots$$

$$\prod_{j=1}^{d_l} p(y_{r_i+j} | s_i) p(D_l \geq d_l) p(s_l | s_{l-1}) \prod_{k=1}^{d_l} p(y_{r_{l-1}+k} | s_l) \quad (1)$$

A simple single-pass, greedy algorithm is used for learning the parameters and complexity of the per-pixel DDHMMs as well as computing the locally most likely state assignment under an AIC [1] based objective function. An unoptimized multithreaded C++ implementation, running on a 3.46 GHz Intel i7 processor, achieves real-time performance. Specifically, continuously updating a DDHMM at each pixel for a video sequence containing seventeen hundred frames with resolution  $240 \times 320$  pixels takes an average of 31 milliseconds per frame.

The Swing video sequence shown in Figure 2 shows a mother pushing her daughter on a swing set and eventually, a previously unobserved pedestrian enters and exits the scene. This seemingly innocuous footage contains interesting periodic phenomena that modern change detection algorithms cannot model. The mother's motions are repetitive as she pushes the child with a periodic rhythm. The mother and daughter on the swing set are considered normal, they are using the swing set for the entirety of the video sequence, and the pedestrian is a change of interest. By modeling intensity persistence, the proposed method is able to explicitly model the dynamics of the swinging child and avoid false positive detections. Competing algorithms exhibit significantly higher false positive rates for this sequence.

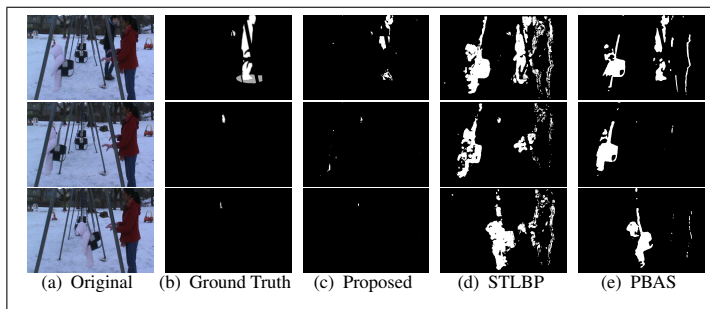


Figure 2: *Swing sequence*: Change detections are visualized as white pixels, normal scene activity as black. The proposed method is the only algorithm which can learn the swinging child is a normal part of the scene but still detect the previously unobserved pedestrian.

The paper discusses the implementation of the online DDHMM learning and inference algorithm as well as the construction of the DDHMM based code book and its application as a classifier for detecting changes in novel video segments. The proposed method is compared to current state of the art change detection algorithms and is shown to be superior, especially in environments containing complex periodic phenomena.

[1] H. Akaike. A new look at the statistical model identification. *Automatic Control, IEEE Transactions on*, 19(6):716–723, Dec 1974. ISSN 0018-9286. doi: 10.1109/TAC.1974.1100705.

[2] Yann Guédon. Estimating hidden semi-markov chains from discrete sequences. *Journal of Computational and Graphical Statistics*, 12(3): pp. 604–639, 2003. ISSN 10618600. URL <http://www.jstor.org/stable/1391041>.