

Plan-Recognition-Driven Attention Modeling for Visual Recognition

Yantian Zha, Yikang Li, Tianshu Yu, Subbarao Kambhampati, Baoxin Li

Arizona State University

{yantian.zha, yikangli, tianshuy, rao, baoxin.li}@asu.edu

Abstract

Human visual recognition of activities or external agents involves an interplay between high-level plan recognition and low-level perception. Given that, a natural question to ask is: can low-level perception be improved by high-level plan recognition? We formulate the problem of leveraging recognized plans to generate better top-down attention maps (Gazzaniga 2009; Baluch and Itti 2011) to improve the perception performance. We call these top-down attention maps specifically as plan-recognition-driven attention (PRDA) maps. To address this problem, we introduce the Pixel Dynamics Network. Pixel Dynamics Network serves as an observation model, which predicts next states of object points at each pixel location given observation of pixels and pixel-level action feature. This is like internally learning a pixel-level dynamics model. Pixel Dynamics Network is a kind of Convolutional Neural Network (ConvNet), with specially-designed architecture. Therefore, Pixel Dynamics Network could take the advantage of parallel computation of ConvNets, while learning the pixel-level dynamics model. We further prove the equivalence between Pixel Dynamics Network as an observation model, and the belief update in partially observable Markov decision process (POMDP) framework. We evaluate our Pixel Dynamics Network in event recognition tasks. We build an event recognition system, ER-PDN, which takes Pixel Dynamics Network as a subroutine, to recognize events based on observations augmented by plan-recognition-driven attention.

Introduction

For agents that need to accomplish tasks in challenging dynamic and multi-agent environments, it is important to be able to selectively focus on certain interesting visual regions. For example, an intelligent surveillance camera may monitor different people doing different activities in different areas. This task requires high visual loads. Selective attention (Sternberg and Sternberg 2016) allows agents to avoid being overloaded, and make the best use of their intra-agent resources. Thereby, the performance and efficiency is expected to be improved. The attention driven by high-level and task-relevant information is called top-down attention (TDA). TDA in humans generally involves more cognitive processes, consuming mental resources like memory, reasoning,

and planning. Selective attention can be driven by top-down feedback, or simultaneously triggered by physical characteristics, using bottom-up attention (BUA). In the field of computer vision, people have traditionally delved into constructing combined BUA and TDA for images, using hand-crafted features (Navalpakkam and Itti 2006). In recent years, a few works using neural networks to learn combined BUA and TDA models have also emerged, like (Anderson et al. 2018; Girdhar and Ramanan 2017; Kosiorek, Bewley, and Posner 2017). Such models may work for static images or video streams. However, none of these has considered if the TDA could be driven by the recognized plans of other agents in the environment, which should be a natural consideration in multi-agent applications. We name this kind of TDA as plan-recognition-driven attention (PRDA). In our work, we address the problem of generating PRDA maps from video frames and corresponding recognized plans.

PRDA may provide practically useful models of selective attention for agents operating in multi-agent environments. In such scenarios, interesting regions far away from an agent are likely to be connected by a plan. And the recognized plan, which is a sequence of actions that an agent is expected to take in the near future, contain useful long-term information for guiding the TDA generation. In this work, we focus on using recognized plans of multiple agents to help generate a PRDA map. We choose event recognition in video surveillance as the primary domain to illustrate the design and evaluation of the proposed method, since surveillance videos often include multiple agents involving in various group activities occurring in various regions of the field of view.

One of the major challenges for using recognized plans is to collect clean observations. This is traditionally done by human or object recognition and tracking. However, in highly cluttered and dynamic environments, methods like object recognition and tracking would become unreliable. For example, humans and objects in the video could be too small to be reliably recognized. Human and object motion can exacerbate the difficulty. Consequently, the collected observations can be incomplete and noisy. To address this problem, we propose a pixel dynamics network (PDN), which learns to directly predict the next state of an object at the position of a pixel. And the prediction is conditioned on both pixel-observation and pixel-level action. Therefore,

it works like a dynamics model in control theory (Branicky, Borkar, and Mitter 1998) – a model that is able to predict the next state of an object, if given the current state, and an action applied to that object.

Note that the pixel-level action does not directly belong to a plan, because there is no *intention* for object points. That said, our PDN could decompose a meta action from a plan, to a group of pixel-level actions, which we would explain in detail later.

To collect meta action features, we extract the Histogram of Optical Flow (HOF), Histogram of Gradients (HOG), and Histogram of Frame Difference (HOD) features from local video frames (video tubes) where an event happens. These video tubes could be obtained by temporally prolonging a bounding box. Then we do clustering on those features, to learn a set of augmented motion primitives (AMPs). Each AMP serves as a meta action, and describes an activity that happens in a video tube.

Our PDN has a subroutine of converting an AMP to pixel-level actions (we denote them as action conditional filters, or ACFs). Then PDN could take a patch of pixels as an observation, and a pixel-level action, to predict the next state of the object point at the center location of that patch. This procedure also has connections to the belief update of the observation model that can handle uncertain observations (i.e., the one in POMDP framework), that we will elaborate later.

To evaluate the effectiveness of PDN, we implement it in an event recognition system. We compare the performance of the event recognition system with our PDN against a state-of-the-art visual recognizer with combined BUA and TDA (Girdhar and Ramanan 2017).

Related Work

Soft Attention Driven by High-Level Signals

Here, we briefly introduce previous work on using neural networks to generate soft attention maps, driven by high-level (or top-down) signals. The work (Girdhar and Ramanan 2017) proposed an attention pooling layer to replace the normal pooling layer in ConvNets. Their high-level signal comes from class labels. For some other previous works that combine top-down and bottom-up attention, their top-down attention model only imitates how spatial reasoning influences attention. In (Navalpakkam and Itti 2006), they use target object locations and surrounding distractor feature in an image as high-level signals to guide the top-down attention for visual search tasks. In (Kosiorrek, Bewley, and Posner 2017), they use the appearance feature of targets. In contrast, our planning-driven attention model exploits both spatial and temporal reasoning for long-term future, based on the awareness of environment dynamics and recognized plan feature of other agents.

Pixel-wise Classification

The design of PDN is also inspired by pixel-wise classification problems (i.e., semantics segmentation). Semantics segmentation is the vision problem that requires assigning a proper semantic label for each pixel. This could be accomplished by using ConvNets (Long, Shelhamer, and Darrell

2015; Noh, Hong, and Han 2015; Carreira et al. 2012; Chen et al. 2014; Mostajabi, Yadollahpour, and Shakhnarovich 2015). Our PDN, instead, assigns a state offset value for each pixel.

Plan Recognition

One of the main novelties of our work is to use high-level signals (recognized plans) to form better attention maps. Plan recognition is about predicting the future action sequences of an agent, given its observed actions. While traditionally people rely on classic planning approaches with a domain model to do plan recognition (Ramírez and Geffner 2009), recent years have seen an increasing interest in doing plan recognition on a learned and approximated model. Zhuo et al., (Tian, Zhuo, and Kambhampati 2016) treat all plans in a plan library as sentences. Then they use the plan library as the training corpora to train a Word2Vec model. The plan recognition task then becomes to find the action (in the action vocabulary) that is most similar to surrounding actions in a testing plan. The similarity value is provided by the learned Word2Vec. Most recently, (Zha et al. 2018) extends the problem range of aforementioned work to using action distribution sequences as plans to learn a Distr2Vec model to handle observational uncertainties.

Our Model

In this section we introduce our pixel dynamics network, and elaborate how it is used in event recognition.

Augmented Motion Primitives Generation

Now we introduce our way of generating AMPs as the representation of pixel-level actions. First, we extract features of HOG, HOF and HOD for each frame. Secondly, we do a K-Means clustering on the total features of the whole dataset. We set the number of clusters to A . We then assign the \tilde{A} nearest centroids of every frame feature. Consequently, we convert a video sequence to an index distribution, which is based on the frame feature clustering. Each distribution has \tilde{A} AMP indices. We also assign a probability value for each index to approximate the uncertainty of an activity. This probability is computed by using the distance between centroids and the corresponding feature. Lastly, we merge several consecutive and identical index distributions to one distribution. This merging of index distribution can make plan traces less redundant.

Thus, recognizing the plan of a group of agents would be realized by predicting a sequence of AMPs for a specific region. The region could either be generated by using bounding boxes, or using attention logits as in our work. Then we assign each of these regions, an identical number m . This indicates that the m -th group agents are doing m -th events. If there are totally M such regions, we label each image pixel with those M numbers. This gives an extra channel (channel $c = 1$) of masking codes, which serves as a mapping from a recognized plan to the corresponding image area. We also add the second extra channel (channel $c = 0$) of pixel locations. After concatenating the two extra channels with original R, G, and B channels, we obtain a masked image,

denoted as V^μ , as illustrated in Figure 1. Correspondingly, the Figure 2 illustrates the procedure of using a pair of local images (located by masking channel) to extract one AMP as in the red circle.

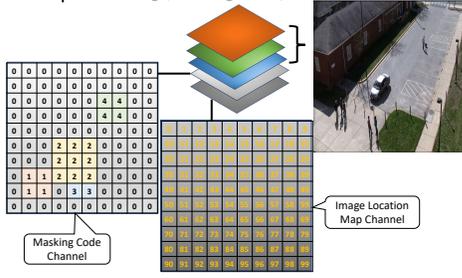


Figure 1: Illustration of V^μ , at one step as a PDN input.

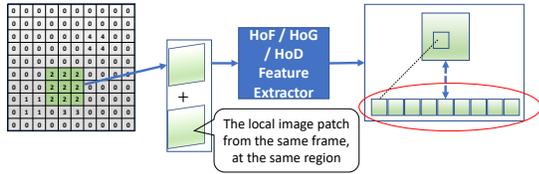


Figure 2: Illustration of extracting one AMP (augmented motion primitive) from a local image patch as a PDN input.

Plan Recognition We adopted the work of Distr2Vec, (Zha et al. 2018), to generate future AMPs, given a sequence of observed AMP distributions. To elaborate, the plan recognition algorithm is based on a shallow planning model, Distr2Vec (distribution to vector), proposed in the work (Zha et al., 2018). Distr2Vec assumes observed action distribution sequences as inputs. This is a good fit for using noisy recognition outputs from a vision model, which consists of sequences of distributions over motion primitives. Distr2Vec then learns an action affinity model from the distribution sequences, by minimizing the distance between each of two action distributions. With this trained action affinity model, plan recognition (predicting future actions) is done, by searching for the most similar action (or motion primitive) to observed action distributions.

Pixel Dynamics Network

Our PDN also takes both a masked image, V^μ (Figure 1), and an AMP (Figure 2), as inputs. With these two inputs, the PDN would generate a state prediction map (SPM). The value at each location of the SPM represents the number of pixels that have been moved to that location by that AMP (action). This is like parallelly applying each of a group of dynamics models to each observation (a local image patch of pixels with corresponding locations) extracted from V^μ . The architecture of PDN is shown in Figure 3. How PDN works could be divided into three stages: action conditional filter (ACF) generation (hidden output $h^{l=3}$ after third layer), convolution between ACFs and the input image to obtain state prediction values in a set of feature maps ($h^{l=4}$ after

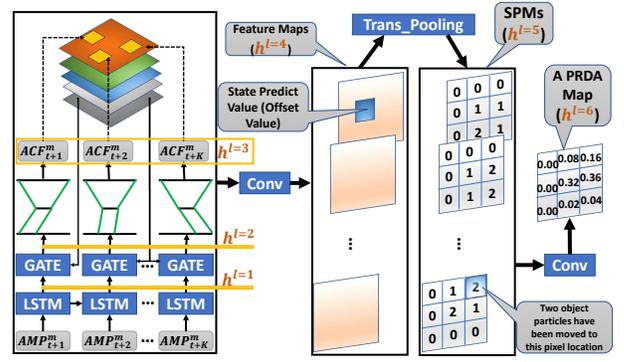


Figure 3: The architecture of Pixel Dynamics Network.

fourth layer), and a translation pooling applied to these feature maps, to get a SPM ($h^{l=5}$ after fifth layer). We will introduce the three stages one by one in the following subsections.

Generation of Action Conditional Filter Because we apply our PDN to generate attention maps driven by recognized plans, we feed the input action to a Long Short-term Memory (LSTM) layer. The LSTM layer could handle the sequential correlations among actions in a recognized plan. Note that an AMP contains pixel-level motion information in a local region (of one event). Hence, we cannot directly use the output of (the first) LSTM layer ($h^{l=1}$) as the filter to convolve with a local image patch. The local image patch being convolved might be smaller than a bounding box or a region with high attention logits. As a result, that local image patch only match a sub-region of the input AMP. This is why we have a gate layer. The gate layer is used to select a certain part of $h^{l=1}$ that matches a local image patch, which leads to the need of inputting the channel $c = 0$, the pixel location map, to the gate layer. The output of (the second) gate layer, $h^{l=2}$, could be seen as a sparse version of the $h^{l=1}$. After applying the gate function, some parts of $h^{l=2}$ are not usable anymore. When doing convolution between $h^{l=2}$ and a local image patch, some parts in a local image patch would not contribute to the computation. The important information in a local image patch may be missed, which could be a problem. We address this problem by proposing a biased encoding-decoding layer that perform encoding-decoding on $h^{l=2}$, which reflects parts of $h^{l=1}$ selected by the gate layer. Thus, it does encoding (compressing) based on input part that is being paid attention to, and reconstruct with richer details (like zooming in). The reconstruction output per step is an ACF. Generating an ACF is to convert a meta action (AMP) to a group of pixel-level actions. $h^{l=3}$ consists of $M \times K$ ACFs. The equation for using the above procedure to compute an ACF, is shown in Equation 1.

$$ACF_{k,t}^m = bias - decoding(bias - encoding(LSTM(AMP_{1:k,t}^m) \odot \sigma(W^{GX} V_t^{\mu,c=0} + B^G))), \quad 1 \leq k \leq K, 1 \leq m \leq M \quad (1)$$

where $ACF_{k,t}^m$ is the k -th action-conditioned filter generated from the k -th AMP in a plan of length K . That plan is the m -th recognized plan, of totally M plans. This is also the plan of m -th group of people that are doing m -th event. We further assume that all of the M recognized plans equally have length K . $LSTM(\cdot)$ denotes that the observed k AMPs are fed into a LSTM cell one by one, and output a LSTM state ($h^{l=1}$) per step. $V_t^{\mu, c=0}$ is the first channel of input image, containing a map of pixel locations. The values in $V_t^{\mu, c=0}$ would range from zero to the product of height of width of input image. The W^{GX} and B^G denote weights and bias parameters of the gate layer. The \odot between LSTM output and gate layer output denotes the element-wise multiplication. The result of this multiplication is fed into a bias-encoding and bias-decoding module.

Our bias-encoding and bias-decoding module works by performing a K-max pooling on the output from gate layer, which obtain a condensed vector, as shown in the small interval (above the gate module) between two green lines in Figure. 3. Then take that condensed input from K-max pooling as the hidden feature, and adopt normal auto-encoder framework to reconstruct the hidden vector to a feature, which is an ACF.

Generation of State Prediction Values With $K \times M$ generated ACFs, the next task is to use each of them, to convolve with input image. Values in output feature maps after the convolution between an ACF and the input image, would be the offset values of the object particle at the center location of local image patch being convolved with an ACF. The convolution between an ACF and the input image could be described by the Equation. 2.

$$h_{ij}^{l=4} = (V_t^\mu \odot W^F * F_{k,t}^m)_{ij} = \sum_{a=0}^{F_1-1} \sum_{b=0}^{F_2-1} [[V_{i+a, j+b}^{\mu, c=0} = m]] \sum_{c=2}^4 (ACF_{a,b}^{m,k} \cdot (V_{c, i+a, j+b}^\mu \cdot W_{i+a, j+b}^{l=4}) + B_{i,j}^{l=4}) \quad (2)$$

where $h_{ij}^{l=4}$ is the output from convolving ACFs with V_t^μ , i.e., the computation of fourth layer. $W_{i+a, j+b}^{l=4}$ and $B_{i,j}^{l=4}$ are weights and bias in the fourth layer, that is used for the convolution with value at location (i, j) . F_1 and F_2 are respectively the height and width of each ACF. $ACF_{a,b}^{m,k}$ is the value at location (a, b) in the k -th ACF of m -th plan. As mentioned earlier, V_t^μ is an masked image at time t , with five channels: pixel location map in channel 0, masking code map in channel 1, and RGB maps from channel 2 to 4. Particularly, the masking code map has values that match the index of a recognized plan. In a frame of image, if there are multiple groups of people doing multiple events, then each group matches a highlighted region, and a recognized plan. All pixels in that region are masked by the index of the recognized plan. The values of masking code are in the masking map. If there are M_t different recognized plans at time t (as shown in Figure. 4, there would be $M_t + 1$ different values in masking map. The reason is, we assign the value zero, to the regions of pixels do not belong to any of recognized plans.

Translation Pooling Once obtaining a feature map of state prediction values, the next step is to convert that to a SPM, which is accomplished by our translation pooling layer. A state prediction value at a position represents the offset as explained before, and the translation pooling layer works by converting that offset value, to a "plus one" operation at a new position, after applying the offset value to the old position. The algorithm of translation pooling is shown in Algorithm. 1.

Algorithm 1 Translation Pooling

```

1: procedure TRANS-POOLING( $h^{l=4}, \text{shape}(N, N)$ )  $\triangleright$ 
   From a set of  $N \times N$  feature maps, in which values are
   computed using Equ. 2
2:    $h_{0:N \times N-1}^{l=5} \leftarrow 0$   $\triangleright$  Initialize all values in the vector
    $h^{l=5}$  (size is  $N^2$ ) to zeros
3:   for  $i = 1; i \leq N; i++$  do
4:     for  $j = 1; j \leq N; j++$  do
5:        $h_{ID(i,j)-h_{ij}^{l=4}}^{l=5} += 1$   $\triangleright$  Generate statistics
   using all state prediction values in  $h^{l=4}$ 
6:     end for
7:   end for
8:    $SPM \leftarrow \text{reshape } h^{l=5} : \text{from } (N^2) \text{ to } (N, N)$ 
9:   return  $SPM$   $\triangleright$  Return a state prediction map
10: end procedure
11: procedure ID( $i, j$ )  $\triangleright$  Convert 2D index to 1D index
12:   return  $i * N + j$ 
13: end procedure

```

PRDA Generation Layer

The role of this layer is to pixel-wise summarization across all channels of SPMs (feature maps) from the translation pooling layer. This summarization will be taken as a PRDA map. Intuitively, as each of those SPMs maps to an action of an agent, summarization of SPMs would provide useful information of which visual region is likely to get considerable amount of pixel objects. As mentioned before, a pixel object means the object point at a location with the corresponding pixel as the observation. The PRDA generation layer also needs to output a feature map that has the same size of a SPM and V_t^μ . We design a special ConvNet layer with filter size, weight, and slide stride set to one. And the weights in this layer are not trainable. This way, we are able to do summarization over all pixel regions, and across all SPMs, which is identical to Equation 3.

$$h_{ij}^{l=6} = PRDA_{i,j}^{t+1} = \frac{\sum_{c=0}^{M \times K - 1} SPM_c^{t,i,j}}{Z} \quad (3)$$

where Z is a normalization number, and we set it to one in our experiments. In $h_{l=5}$, there are $M \times K$ SPMs, because there are M plans of length K , and each AMP leads to a SPM. Each SPM has $N \times N$ values. Generating a PRDA map requires calculating statistics over all of them. In addition, here c denotes the c -th channel or SPM in $h_{l=5}$, t denotes the current time step, and each pixel location in a PRDA map is identified by i and j .

Event Recognition with PDN

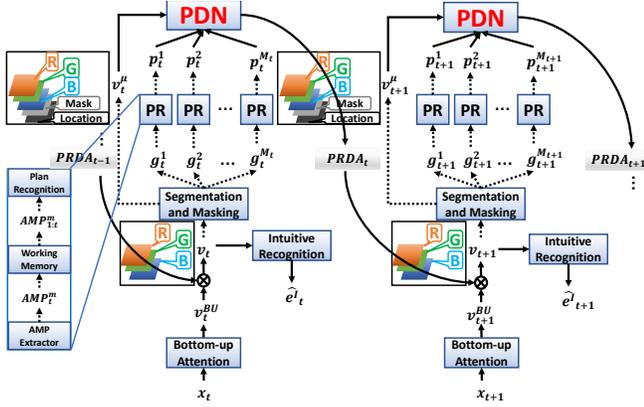


Figure 4: The event recognition system with Pixel Dynamics Network.

Here we show how our PDN can be used to generate PRDA maps to support an event recognition task. We denote our event recognition system as ER-PDN. The overall framework is shown in Figure. 4. The solid paths are differentiable, and dashed ones are not. At each step, a video frame X_t , is filtered by bottom-up attention model, to get a BUA map V_t^{BU} . The V_t^{BU} is then combined with a PRDA map from the previous step, to get a highlighted frame V_t . The segmentation and masking component then extracts M_t glimpses, which are disconnected highlighted areas in V_t . Each of these glimpses are fed into a separate but structurally identical plan recognition component, PR, to generate M_t recognized plans p . The segmentation and masking component also adds additional channels as explained before, to map each of M_t recognized plans that is indexed by m , to a highlighted region in V_t^μ . Consequently, every pixel in V_t^μ also has a mask code m . The PDN takes all recognized plans for different highlighted areas, and the masked V_t^μ , to generate the PRDA map $PRDA_t$.

The above procedure describes the non-differentiable path (the dashed ones). In order to train the model, we want to let the model to predict a label so that there is a teaching signal. This is what the dashed path is showing. After generating V_t^μ , an intuitive recognition module is required to predict an event label at each step. The intuitive recognition module is a LSTM layer, which has the advantage of handling sequential inputs with internal long-term dependency.

Deeper Analysis of PDN

Connection to Belief Updates in Observation Models

We start from providing the belief update procedure in the observation model, in partially observable Markov decision process (POMDP) framework (Monahan 1982; Littman 2009). Observation model could handle observation uncertainties, and model the evolution of states. POMDP framework assumes state uncertainty, and treats actions as triggers between states (the essential difference to belief updates in

other observation models, e.g., HMMs). The belief update process in POMDP is described in Equation 4, 5, 6, and 7. In Equation 4, $b_0(s)$ denotes the initial belief of state s . Equation 5 and 6 define the observation model and transition model respectively. Equation 7 defines how belief is updated.

$$b_0(s) = Pr(S = s) \quad (4)$$

$$O(s', a, o) = Pr(o|s', a) \quad (5)$$

$$T(s, a, s') = Pr(s'|s, a) \quad (6)$$

$$b'(s') \propto O(s', a, o) \sum_{s \in S} T(s, a, s') b(s) \quad (7)$$

Now we elaborate the connection between our PDN and POMDP's observation model. PDN takes an image and an action (later becomes to an ACF) as inputs, and predict how each object point at each image location (each location maps to a pixel value) would be shifted. The state s denotes which object point at which image location. We assume this state is hidden. We can only guess this state from the observation of a group of pixels. Thus the observation uncertainty in PDN essentially means: given the observation of a group of pixels (e.g., a square patch of pixels), the true location of the object point that should be in the center of that patch, could actually be at anywhere in that patch.

Equation 2, which predicts location offset values of a certain object point ($h_{i,j}^{l=4}$), works like using a ConvNet to model the belief updates (Equation 7). Equation 5 and 6 are subroutines of Equation 7. To predict the s' in Equation 6, the neural network that approximates $T(s, a, s')$ needs the inputs of an image location with the corresponding pixel value, and the corresponding value in an ACF. If we multiply a value from an ACF with a pixel value in the image location (i, j) , we are assuming (or believing) that that pixel is the proper representation of the object point at location (i, j) . Thus $T(s, a, s')$ in POMDP matches the entry of $\sum_{c=2}^4 (ACF_{a,b}^{m,k} \cdot (V_{c,i+a,j+b}^\mu \cdot W_{i+a,j+b}^{l=4}) + B_{i,j}^{l=4})$ in Equation 2 (state s matches $V_{i+a,j+b}^c$, and action a matches $ACF_{a,b}^{m,k}$). The summarization over states, $\sum_{s \in S}$, in Equation 7 (belief updates in POMDP), matches $\sum_{a=0}^{F_1-1} \sum_{b=0}^{F_2-1}$ in Equation 2. These three matches intuitively mean that: the PDN believes that all image locations in a local image patch (to be convolved with an ACF), and the corresponding pixel value could represent the true object point state with certain probabilities. Also, note that when PDN predicts the offset value, PDN is not really moving that object point, but is only observing. Hence, there is no uncertainty in whether an object point is really moved to the target image location or not. In other words, $T(s, a, s')$ is deterministic, and $O(s', a, o) = 1$ (for Equation 5). If we further assume that the belief of states obeys to uniform distribution, i.e., all locations in a local image patch are equally possible to be where the main object point is, then $b'(s') \propto \sum_{s \in S} T(s, a, s')$, which can be reduced to solving the Equation 2.

Relationship with Other Models that can do Pixel Transformation

Spatial Transformer Networks (STN) In the work of STN (Jaderberg et al. 2015), a localization network figures out a set of proper transformation parameters, with the input of an image. These parameters are fed into a grad generator and sampler to transform feature maps. Thus, pixel-level transformation leads to image-level transformation, like rotation, crop, etc. Our proposed PDN shares a key characteristics with STN, with respect to learning to manipulate pixels (pixel transformations). That being said, STNs pixel transformation has a different purpose than PDN, which leads to different designs of model architectures and algorithms. STN intends to warp an entire image, which requires that all pixels in new feature maps still maintain proper values such that they all together form a transformed image. However, for PDN, the purpose is to figure out which image position has been moved a comparable number of object points given pixel-level actions. This statistics indicates proper attention importance. The hidden layer for generating SPMs treats pixels with their corresponding image locations, as observations, and object point with a location, as states. That hidden layer generates a set of state offset values, each reflects where the next image position of an object point is believed to be (belief update).

Dynamic Filter Networks The design philosophy of DFNs (Jia et al. 2016) is coincidentally similar to the ACF generator in PDN: instead of learning filter weights in ConvNets, they have a neural network model, trained to generate values of filter weights on the fly. Thus, DFNs can be seen as a more general version of ACF generator. And the ACF is specifically used to generate tuned pixel-level action features (AMPs), for each image region.

Evaluation

In this section, we provide experiment results and analysis. We compare our PDN with a baseline vision model (Girdhar and Ramanan 2017), by using them to solve an event recognition task. Note that the baseline model also has an attention layer. Unlike existing event recognition works which only use event-relevant sequences (Wang and Ji 2015; Mahmud, Hasan, and Roy-Chowdhury 2017), we did not do pre-segmentation or use bounding boxes. We ran our experiments on a machine with a Quad-Core CPU (Intel Xeon 3.4GHz), a 64GB RAM, a GeForce GTX 1080 GPU, and Ubuntu 16.04 OS.

Dataset Pre-processing and Training Procedure

We used the VIRAT 1.0 Ground Dataset (Oh et al. 2011) for evaluation. We randomly select 70% of all videos for training, and the rest for testing. And we down-sampled videos to use every six frames. We used different pre-processing procedures for learning augmented motion primitives (AMPs), and learning the PDN.

To learn motion primitive clusters, we applied K-means clustering to find the motion primitive cluster library. The motion primitive cluster library is a library of mappings between a motion primitive vector, and a cluster index. We set

the number of clusters to 512, and apply clustering on HoD features collected from 70% of VIRAT videos for training. We did not use the original VIRAT videos. Instead, we used the video segments that are inside ground-truth bounding boxes. This means that the videos we use are both temporally and spatially smaller than the original videos. Those bounding boxes highlight event relevant frames and objects. In contrast, when training the ER-PDN, we did not use any bounding box information in videos, which significantly increases the event recognition difficulty. Once we obtain the AMP clusters, we can convert sequences of HoD features to a sequence of motion primitive (or AMP) indices. Each of such a motion primitive index sequences could be seen as a high-level plan. Then we use the work of (Zha et al. 2018) to learn a shallow plan recognition model. When training the plan recognition model, we set the distribution size to three, and training epoches to sixty.

When training ER-PDN and the baseline model (Girdhar and Ramanan 2017), We feed both models with resized images (450×450). Note that if there are multiple events simultaneously happening, then some video frames may map to multiple labels. For this situation, we copy certain images multiple times, and assign each with a corresponding event label. The video down-sampling allows us to increase the variance between each of two frames in our pre-processed dataset, and thereby increase the variance between each of two AMPs in a recognized plan.

When we train the ER-PDN, we feed four frames as a batch to the pre-trained plan recognition module to generate a K -step recognized plan (we set K to five). With two consecutive frames, the HOD feature extraction component generates a HoD feature vector, which would be used for searching top-3 AMP indices from the trained AMP cluster library. Thus, we could obtain a sequence of three AMP index distributions (each distribution has three AMPs), given an input image batch. This is similar to the work (Zha et al. 2018). This AMP distribution sequence is used as an observed plan, fed into the plan recognition module, which returns a k -step sequence of future AMP indices that is most likely to happen (the recognized plan). ER-PDN then converts the future AMP indices back to a sequence of motion primitive feature vectors, by using the AMP cluster library. ER-PDN then feeds these vectors one by one, to generate proper ACFs, and SPMs. This is how ER-PDN uses input actions (AMP vectors).

To generate a PRDA, ER-PDN also takes an image as input. After taking a batch of four consecutive images from a video, ER-PDN uses the first image to generate a BUA, which is then used to generate glimpse for plan recognition. And ER-PDN uses the last image, with the generated PRDA and BUA map, to augment the input and predict which event is happening.

Results and Analysis of Event Recognition

A comparison of the results in Table 1 shows that the visual recognition could be improved by a plan (a sequence of motion primitives). In addition, Figure 5 contains examples of input video frames, and the corresponding feature maps of BUA, PRDA, and V_t^μ . Specifically, in Figure 5, (a), (b) and

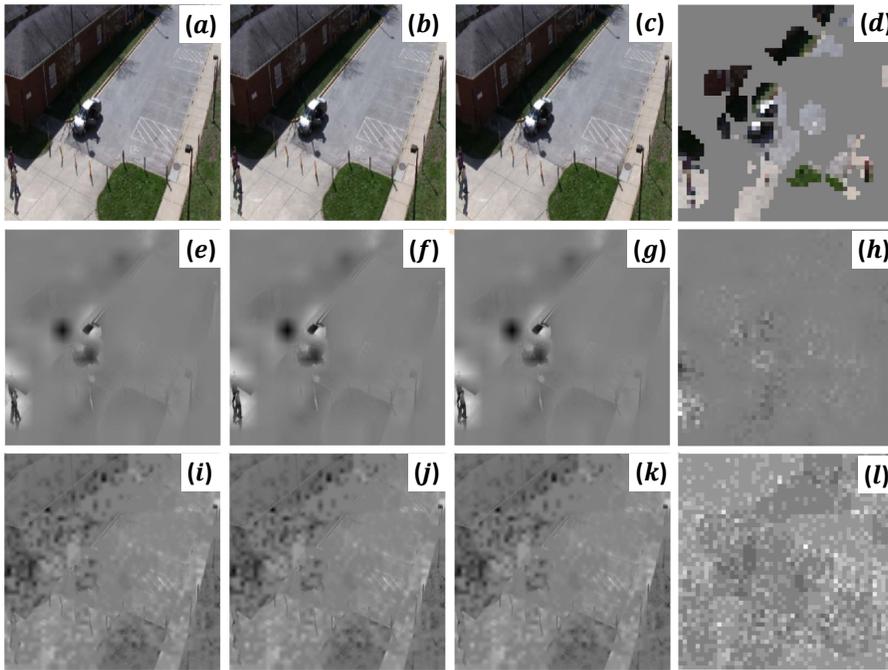


Figure 5: Results Visualization of our Event Recognition System with the PDN.

Model	Mean Average Precision
Baseline (Girdhar and Ramanan 2017)	29.06%
ER-PDN	29.15%

Table 1: Statistics of event recognition performance on VIRAT dataset, using the baseline model (Girdhar and Ramanan 2017), and our ER-PDN.

(c) are resized video frames. (d) is a visualization of V_t^μ , by applying a BUA to a resized frame and keeping parts within a threshold. (e), (f), and (g) are visualizations of applying a BUA to a resized frame. (h) is the feature map of applying both a BUA and a PRDA to the output of a convolutional layer, with an input frame. (i), (j) and (k) are visualizations of applying a PRDA to input video frames. (l) is a PRDA feature map.

We can observe from (a), (b) and (c) in Figure 5 that, three people are moving toward a car. (e), (f), and (g) in Figure 5 indicate that, the BUA map focus properly to interesting image areas, where there is a car, and people. Because the car is not moving, in (i), (j) and (k), the area with car is darker, which indicates less attention. The area between the three people, and the car, is lighter, which suggests that the moving people may want to get to the car. Thus, the lighter region between the people and the car indicates that there is more attention put there.

Conclusion and Future Work

In this work, we formulated the problem of generating attention maps which are driven by plans. While we assume that the plans come from a plan recognition module, the plans

could also come from, e.g., an agent itself. Essentially, those plans contain long-term information that an agent could make use of, in order to focus on proper regions that are closely relevant to a task.

To address this problem, we proposed a Pixel Dynamics Network (PDN) that learns a dynamics model for object points in video frames. We assume that, in each pixel location, there is an object point, and PDN learns to predict the next pixel location (as the next state) of that object point, on an observation of several pixels, and a motion primitive vector (as an action). Furthermore, if we feed PDN with a sequence of motion primitives (i.e., a plan), then PDN could predict a sequence of future states. If we apply PDN to all object points that a video frame could reflect, and all actions from one or more plans, we could obtain a set of feature maps, that show a statistics of K -step future states of all object points. This set of feature maps are called state prediction maps (SPMs). By summarizing them, we could obtain the plan-recognition-driven-attention map (PRDA). We empirically show that an event recognition task could be improved by using PRDA maps, i.e., the plan-recognition-driven-attention modeling.

In future work, we plan to leverage not only HoD features, but HoF, and HoG feature as well, to learn better motion primitive features, to improve PDN performance.

Acknowledgements

This research is supported in part by the AFOSR grant FA9550-18-1-0067, the ONR grants N00014161-2892, N00014-13-1-0176, N00014-13-1-0519, N00014-15-1-2027, and the NASA grant NNX17AD06G.

References

- Anderson, P.; He, X.; Buehler, C.; Teney, D.; Johnson, M.; Gould, S.; and Zhang, L. 2018. Bottom-up and top-down attention for image captioning and visual question answering. In *CVPR*, volume 3, 6.
- Baluch, F., and Itti, L. 2011. Mechanisms of top-down attention. *Trends in neurosciences* 34(4):210–224.
- Branicky, M. S.; Borkar, V. S.; and Mitter, S. K. 1998. A unified framework for hybrid control: Model and optimal control theory. *IEEE transactions on automatic control* 43(1):31–45.
- Carreira, J.; Caseiro, R.; Batista, J.; and Sminchisescu, C. 2012. Semantic segmentation with second-order pooling. In *European Conference on Computer Vision*, 430–443. Springer.
- Chen, L.-C.; Papandreou, G.; Kokkinos, I.; Murphy, K.; and Yuille, A. L. 2014. Semantic image segmentation with deep convolutional nets and fully connected crfs. *arXiv preprint arXiv:1412.7062*.
- Gazzaniga, M. S. 2009. *The cognitive neurosciences*. MIT press.
- Girdhar, R., and Ramanan, D. 2017. Attentional pooling for action recognition. In *Advances in Neural Information Processing Systems*, 34–45.
- Jaderberg, M.; Simonyan, K.; Zisserman, A.; et al. 2015. Spatial transformer networks. In *Advances in neural information processing systems*, 2017–2025.
- Jia, X.; De Brabandere, B.; Tuytelaars, T.; and Gool, L. V. 2016. Dynamic filter networks. In *Advances in Neural Information Processing Systems*, 667–675.
- Kosiorrek, A.; Bewley, A.; and Posner, I. 2017. Hierarchical attentive recurrent tracking. In *Advances in Neural Information Processing Systems*, 3053–3061.
- Littman, M. L. 2009. A tutorial on partially observable markov decision processes. *Journal of Mathematical Psychology* 53(3):119–125.
- Long, J.; Shelhamer, E.; and Darrell, T. 2015. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 3431–3440.
- Mahmud, T.; Hasan, M.; and Roy-Chowdhury, A. K. 2017. Joint prediction of activity labels and starting times in untrimmed videos. In *Computer Vision (ICCV), 2017 IEEE International Conference on*, 5784–5793. IEEE.
- Monahan, G. E. 1982. State of the art survey of partially observable markov decision processes: theory, models, and algorithms. *Management Science* 28(1):1–16.
- Mostajabi, M.; Yadollahpour, P.; and Shakhnarovich, G. 2015. Feedforward semantic segmentation with zoom-out features. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Navalpakkam, V., and Itti, L. 2006. An integrated model of top-down and bottom-up attention for optimizing detection speed. In *null*, 2049–2056. iee.
- Noh, H.; Hong, S.; and Han, B. 2015. Learning deconvolution network for semantic segmentation. In *Proceedings of the IEEE international conference on computer vision*, 1520–1528.
- Oh, S.; Hoogs, A.; Perera, A.; Cuntoor, N.; Chen, C.-C.; Lee, J. T.; Mukherjee, S.; Aggarwal, J.; Lee, H.; Davis, L.; et al. 2011. A large-scale benchmark dataset for event recognition in surveillance video. In *Computer vision and pattern recognition (CVPR), 2011 IEEE conference on*, 3153–3160. IEEE.
- Ramírez, M., and Geffner, H. 2009. Plan recognition as planning. In *IJCAI 2009, Proceedings of the 21st International Joint Conference on Artificial Intelligence, Pasadena, California, USA, July 11-17, 2009*, 1778–1783.
- Sternberg, R. J., and Sternberg, K. 2016. *Cognitive psychology*. Nelson Education.
- Tian, X.; Zhuo, H. H.; and Kambhampati, S. 2016. Discovering underlying plans based on distributed representations of actions. In *Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems*, 1135–1143. International Foundation for Autonomous Agents and Multiagent Systems.
- Wang, X., and Ji, Q. 2015. Video event recognition with deep hierarchical context model. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 4418–4427.
- Zha, Y.; Li, Y.; Gopalakrishnan, S.; Li, B.; and Kambhampati, S. 2018. Recognizing plans by learning embeddings from observed action distributions. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems*, 2153–2155. International Foundation for Autonomous Agents and Multiagent Systems.