Review of Video Predictive Understanding: Early Action Recognition and Future Action Prediction

He Zhao and Richard P. Wildes

Lassonde School of Engineering

Department of Electrical Engineering and Computer Science

York University

Abstract

Video predictive understanding encompasses a wide range of efforts that are concerned with the anticipation of the unobserved future from the current as well as historical video observations. Action prediction is a major sub-area of video predictive understanding and is the focus of this review. This sub-area has two major subdivisions: early action recognition and future action prediction. Early action recognition is concerned with recognizing an ongoing action as soon as possible. Future action prediction is concerned with the anticipation of actions that follow those previously observed. In either case, the **causal** relationship between the past, current and potential future information is the main focus. Various mathematical tools such as Markov Chains, Gaussian Processes, Auto-Regressive modeling and Bayesian recursive filtering are widely adopted jointly with computer vision techniques for these two tasks. However, these approaches face challenges such as the curse of dimensionality, poor generalization and constraints from domain specific knowledge. Recently, structures that rely on deep convolutional neural networks and recurrent neural networks have been extensively proposed for improving performance of existing vision tasks, in general, and action prediction tasks, in particular. However, they have their own shortcomings, e.g., reliance on massive training data and lack of strong theoretical underpinnings. In this survey, we start by introducing the major sub-areas of the broad area of video predictive understanding, which recently have received intensive attention and proven to have practical value. Next, a thorough review of various early action recognition and future action prediction algorithms are provided with suitably organized divisions. Finally, we conclude our discussion with future research directions.

Table of Contents

List of Figures

List of Tables

Chapter 1

Introduction

1.1 Motivation

In the computer vision community, video predictive understanding refers to the inference of the un-observed future from the current as well as historical information. From a theoretical perspective, video predictive understanding is of importance as it addresses the basic scientific question of how to extrapolate future events from previous observations. From a practical perspective, the ability to predict future events from videos has potential to provide positive impact for wide range of artificial intelligence applications, including autonomous vehicle guidance, effective time response to surveillance data and improved human computer interaction.

In video understanding, researchers are generally pursuing answers for two fundamental focal points: **what to predict** and **how to predict**. The intuition behind the first focus is that the future can be of several forms (i.e., future semantic, physical trajectories, poses or pixel distributions) and it is non-trivial to explore its most feasible format to better assist real-life applications. To this end, diverse sub-areas have appeared that work on predicting the future in various knowledge domains. The second focus explores various mathematical techniques, information sources and data modalities to enhance the prediction accuracy.

1.2 Overview of video predictive understanding

The major sub-areas of video predictive understanding are pixel prediction, semantic segmentation prediction, trajectory prediction, human motion prediction, predictive coding as well as early action prediction and future action prediction. While the focus of this review is early action prediction and

1

future action prediction, in the next several paragraphs, each of these sub-areas will be briefly presented to provide context.

Video pixel prediction is concerned with generating pixelwise values for potential future images extrapolated temporally from previous observations. Early approaches to this problem attempted generation through learning procedures without adequate attention to basic elements of image structure (e.g., spatial gradients and optical flow) and yielded results that had undesirable artifacts (e.g., blurry images) as well as short prediction horizons [112, 119, 146]. Subsequent work that explicitly considered such basic image structure yielded improved results [96, 104]; however, prediction horizons were little extended. To support longer term predictions, work has relied on higher-level image analysis (e.g., body keypoint and motion analysis [114, 158, 159]). Other work has made use of stochastic generative models to capture uncertainty [5, 84, 157]. Yet other work has decomposed motion and appearance for separate processing via latent variables to better control outcomes [50, 54, 128, 156] as well as emphasized object centric motions [168, 172, 173]. Generation of naturalistic future images beyond short time horizons remains challenging. Also, it is worth noting that such predictions may lend little insight into higher-level (e.g., semantic) understanding of future events.
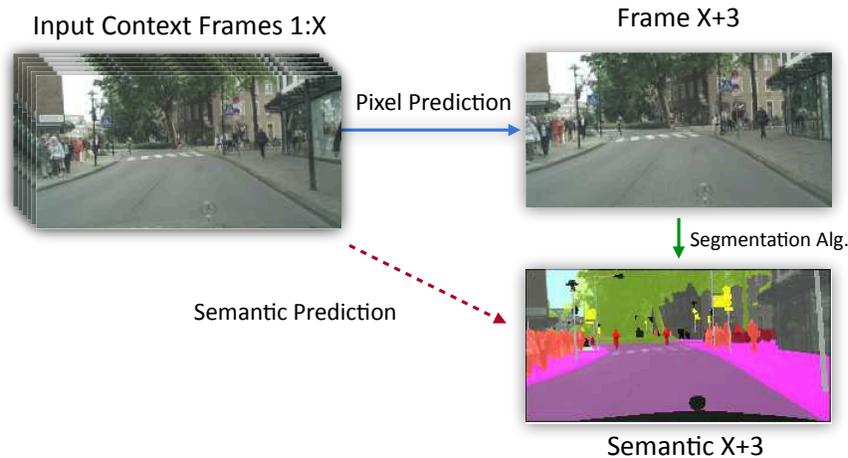
Figure 1.1: Video pixel prediction in upper row and semantic segmentation prediction in lower row. A desired future scene parsing (semantic segmentation of a future frame) can be obtained via either: **1.** pixel prediction and then semantic segmentation; **2.** direct semantic prediction.

Video segmentation prediction bears high similarity with pixel prediction, yet it differs in that its outputs come in terms of pixelwise semantic labels (e.g., road, tree, etc.). To achieve future semantic segmentations, one might either perform pixel prediction followed by standard semantic segmentation operations or directly infer future semantic segmentations. The latter approach is the focus of video semantic segmentation prediction. Figure 1.1 provides a visual illustration of the differences between pixel and semantic segmentation prediction. While different in output (i.e., single valued pixel labels vs. RGB

values), semantic segmentation work has largely leveraged similar techniques as used for pixel predictions, with various levels of emphasis on purely appearance vs. dynamics semantics [58, 97, 98, 140, 171]. An important distinction is that the objective functions for semantic segmentation approaches typically rely on the downstream task (e.g., instance segmentation) compared to pixel predictions.

Video trajectory prediction is another intensively investigated sub-area of video predictive understanding. In this case, observed video is preprocessed into sequences of 2D coordinates that capture the trajectories of objects of interest (e.g., humans, vehicles). The goal is to predict a future sequence in the same format. Early work focused on the forecasting of a single trajectory [63], whereas more recent research has expanded to the prediction of crowd trajectories [3, 40, 53, 91, 134, 134, 135]. Among these efforts, some implicitly fused crowd information, by passing messages between targets, to better infer every single future trajectory [3, 40]. Other work actively embraced rich scene context to assist their predictions [91, 134]. Yet other work explicitly used physical social distance to encode the crowd effect for their trajectory forecasting [53, 135]. Recently, a spatio-temporal version of a graph neural network has been found useful for this task [108]. As with other areas of video prediction, long term prediction remains an outstanding challenge and while some work has made use of scene context, it is likely that much more can be explored in that direction.

Human motion prediction focuses on foreseeing future human motion in the form of body key-points (e.g., 24 body part coordinates [102]). This task is very akin to trajectory prediction, since its data is abstract key-points rather than raw pixels. The majority of work in this sub-area used recurrent neural networks in a Encoder-Decoder fashion to produce future pose sequences [26, 102]. On top of that, a wide range of machine learning techniques have been applied to this sub-area, including attention [153], contrastive learning [17] and imitation learning [161]. An important (and remaining) challenge in this area is to restrict generated motions to be realistic, even while varied. Advances along this direction have exploited geometric constraints [37] as well as generative modeling [4, 161].

Video predictive coding is rather different from the above sub-areas. It is concerned with coding raw video data into a high dimensional representation in such a way that it can be easily decoded into the future of input video. The learned representations are used as video features for various tasks. For example, early work along this line conducted video prediction learning with an objective to extrapolate frame-wise deep features that ultimately were applied to unsupervised action recognition [146]. Follow on work enriched this idea with contrastive learning [44]. A recent effort that tried a similar idea on unsupervised video skeleton action recognition also observed success [147]. Another effort found predictive coding meaningful for video retrievals [177]. A particularly interesting aspect of this line of work is that it requires no annotation, as the learning and feature distillation simply makes use of already available video

data. Researchers find this valuable since predictive coding is self-contained and requires zero annotations, thus it alleviates the heavy burden of human labeling.

Action prediction is principally concerned with the prediction of human actions. Among many important research topics in video understanding, action recognition arguably has received the most interest and efforts in the last decade [59,65]. With tremendous advances achieved and various experiments explored in recent years, some researchers have refocused on the ability to *foresee* the happening of certain actions rather than recognize the completed ones. The reason can be easily seen from the practical application point of view: Full action executions often take elongated time spans and thus performing the recognition algorithm after the action is completed would naturally lose the capability to react in real-time, which, however, in many conditions is essential.

(a)   (b)   (c)
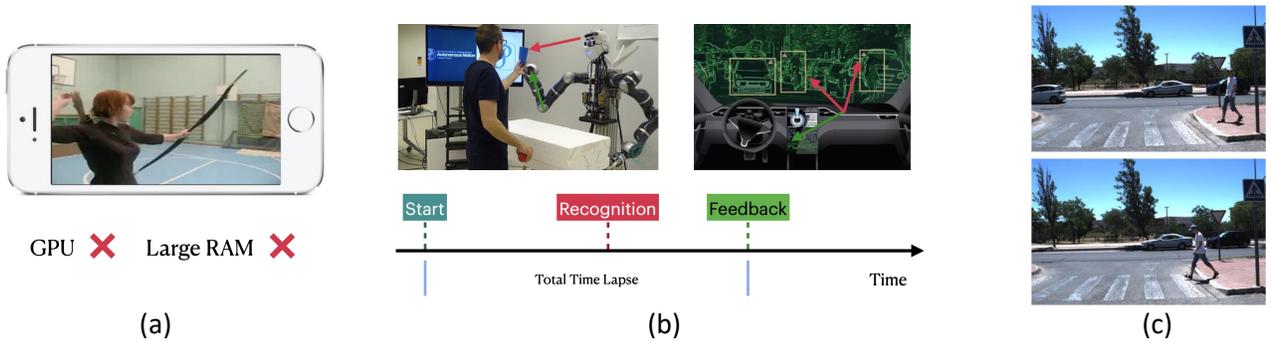
Figure 1.2: Representative applications of early action recognition and future action prediction. Sub-figure (a)-(b) depict the use of early action recognition on reducing the computational budget of portable devices (e.g., mobile phones typically lack GPU and memory space to process elongated videos) and on diminishing the response gap for fast feedback systems (e.g., human-robot interactions and self-driving vehicles). Sub-figure (c) plots the application for future action prediction: Anticipating pedestrian behaviors at traffic intersections.

As examples: In the video anomaly detection task, people expect a vision system that automatically spots any potentially disruptive events as early as possible and then sends out an alarm before any damage is done. Obviously, such a system is required to either recognize harmful actions in effective time or anticipate their happening in advance in space and time. Similarly, in the automonous driving setting, desirable unmanned vehicles need to be highly sensitive to pedestrian walking intentions, in particular at traffic crossings, to avoid accidents. To achieve that goal, the vehicle system needs to leverage contextual information from the scene (e.g., traffic lights and crossings) together with pedestrian ego-motion, standing pose, heading angle, etc., to judge their most likely next action and respond in kind.

In general, based on application scenarios, researchers mainly explore two finer divisions: early action recognition and future action prediction. The former refers to recognizing the on-going actions as early as possible, while the latter to making plausible predictions on future actions that would happen after

the execution of the current one. The focus of this review is on early action recognition and future action prediction. Figure 1.2 provides illustrative examples of both early action recognition and future action prediction.

1.3   Outline

This chapter has served to motivate and provide a broad overview of video predictive understanding. Chapter 2 will provide a detailed review of early action recognition. Chapter 3 will provide a detailed review of future action prediction. Finally, Chapter 4 will provide conclusions, including suggestions for future research directions.

Chapter 2

Early Action Recognition

2.1 Overview

One extensively investigated sub-area of video predictive understanding is early action recognition. Its core mission is to recognizing actions while they are still evolving and predicting the correct action labels as early as possible, as illustrated in Figure 2.1.
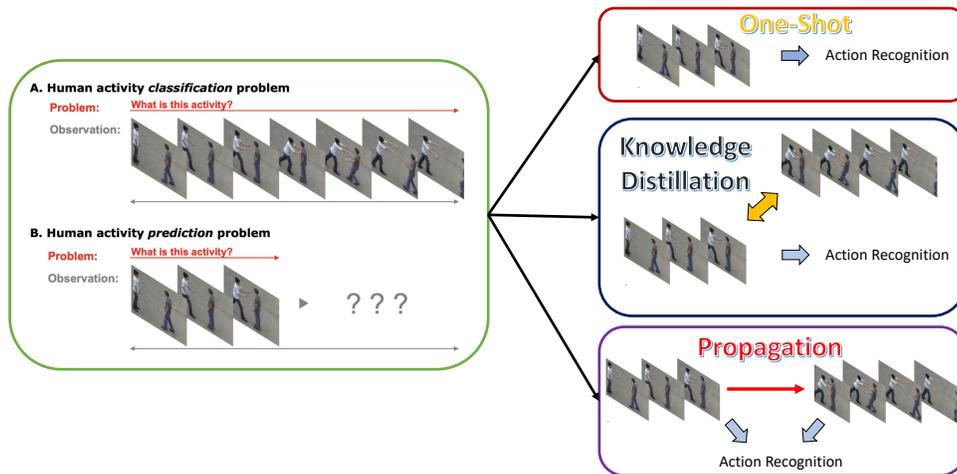
Figure 2.1: (Left) The definition of early action recognition from [129]. Preliminary observations of a particular ongoing action clip are used as input to generate the corresponding class label. Figure modified with permission from [129]. (Right) Illustration of three major perspectives in current literature to tackle early action recognition: One-shot mapping, knowledge distillation and propagation.

Formally, the problem can be defined as: Given a sequence of action video frame observations $y_{1:T}$ and its semantic label $\mathcal{X}$, where 1 and T represent the starting and ending indices of action frames, define a function $f : y_{1:t} \rightarrow \mathcal{X}$, which maps various partial segments $y_{1:t}$ to the groundtruth label domain. In

particular, it is favorable to have high recognition performance on relatively small frame index (i.e., $t \to 0$), meaning that the algorithm is able to recognize actions correctly with very limited initial observations.

Apparently, there is a fundamental assumption behind this problem setting: The full action sequence supports superior recognition compared to observation of only an initial segment, otherwise the problem is unnecessary. To validate this assumption, the original advocate of this research direction, M. S. Ryoo, conducted a series of experiments that operated existing action recognition algorithms on various partial action clip segments (i.e., $y_{1:t}, t \in (1, \ldots, T)$) and demonstrated that insufficient observations in action temporal scope resulted in much inferior recognition accuracy [129]. Also, as a visual example, the partial observation in Figure 2.1 (bottom row) can be sensibly categorized into **pushing** or **hugging** with equal chance, whereas the full action sequence (top row) reveals high likelihood for **pushing**.

Follow on work defined a systematic evaluation procedure for early action recognition [68]: Uniformly divide a full video $y_{1:T}$ into K segments $y[\frac{T}{K}(k-1) + 1, \frac{T}{K}(k)]$ where $k = 1, ..., K$ is the index of the segment. The length of each segment is $\frac{T}{K}$ and different videos usually have separate segment lengths. Thus, a partial observed video sequence can be described as $y[1, \frac{T}{K}(k)]$ and recognition accuracy at all possible k should be reported (e.g., they pre-define $K = 10$ and report 10 sets of accuracies). The majority of successive work follows this paradigm.

In the last decade, researchers approach this task from various perspectives and hereby we group them into three major divisions: **1.** formulating one-shot mappings from partial observations to groundtruth labels of full observations; **2.** distilling the information from the full action clip observations into partial observations; **3.** propagating the limited partial information into the future in a temporal extrapolation fashion. In the following sections, we review all recent research according to these divisions.

2.2 One-Shot Mapping Based

The basic assumption underlying one-shot mapping based approaches is that a partial observation of an action video provides sufficient information in and of itself to define the appropriate overall action class regardless of the unobserved clip portions. What appears to be the first work along these lines [129] made use of sequential feature matching with an Integral Bag-of-Word (IBoW) representation [142]. The intuition being that the feature distributions of partial and complete action clips should be similar, if they derive from the same action. To avoid inefficiencies that would result from training classifiers for each temporal segment, they initially extract Spatio-Temporal Interest Point features (STIP [81]) from raw videos and re-organize them into visual word histograms (IBoW) according to color distributions. Then, a Gaussian kernel matching function is used to generate a probability score representing the likelihood

between partial and whole observations. Finally, the selected class label is that of the full action that produces the largest likelihood. As illustrated in Figure 2.2, for each action segment, a color histogram is calculated. Similar work improved the feature representation as well as selection with sparse coding but shared the overall structure [10].

(a)

(b)
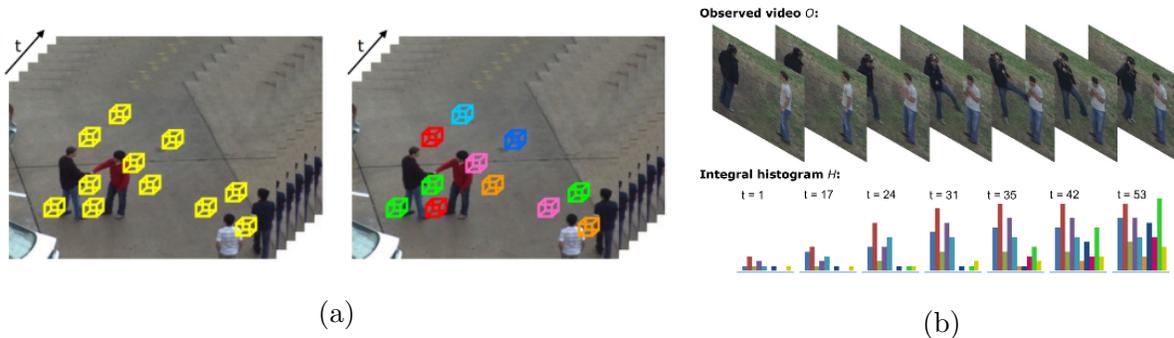
Figure 2.2: (a) STIP features from videos and clustered results by color for each frame. (b) Demonstration of Integral Bag-of-Word (IBOW) used in [129]. For each time step, its feature histogram amasses all previous observations. Figure reproduced with permission from [129].

While the IBoW approach supports fast inference and is learning free, it operates under two restrictive assumptions. First, the feature extractor, STIP, yields an appearance based feature descriptor (indeed STIP comes from spacetime Harris-corner detection [45]), which is sensitive to view changes, illumination and other nuisance details of image acquisition. Thus, targets from the same action category are assumed to be similar in appearance. Second, the scheme to matching feature histograms between partial and full actions is not always practical since actions can evolve dramatically, to the extent that the beginning appears totally unlike the ending.

Subsequent work following essentially the same framework improved on the original by employing stronger features as well as more sophisticated machine learning-based classifiers [80]. More specifically, the feature representation was enhanced by merging motion descriptors (iDT [163]) with appearance descriptors (HOG [18]). Furthermore, to capture the human movements at all levels, they proposed an approach to capture features at three hierarchical temporal scales: a coarse level feature that processes generic videos by taking frames as a whole; a mid-level feature that captures viewpoint specific information by extracting features only belonging to certain specific actors (e.g., *level2*, focuses on regions that are pre-clustered by different human actor appearances); a fine-grained movement level that captures pose specific information (e.g., *level3* extracts features only from regions that share the same actor appearance *and* pose type, which is again obtained by clustering). To finally assign class labels to given video clips, the authors adopted Support Vector Machines (SVMs) [16] that were trained on features from all three levels to jointly make the decision.

Though this relatively new and detailed approach advanced its predecessors through adopting stronger features, hierarchical extractions and learning-based classifiers, it still relied on the assumption that fine grained analysis of an initial partial observation can bridge the past and future. A notable downside to this particular extension is that it relied on clustering algorithms for grouping actors and poses [80], which can be noisy (e.g., it is sensitive to the color of clothing).

Subsequent work [141] disposed of unsupervised clustering for detection of actors and their poses by instead making use of object/human detectors pretrained on massive datasets (e.g., SSD [95]). Therefore, they proposed to jointly perform action localization as well as action early recognition, based on the assumption that actions only associate with behaviors of human actors, and called the detected actor regions **action tubes** (i.e., bounding box coordinates of detected actors). For given frames, this approach generates bounding box regions together with class labels for all possible regions within each frame and rejects negative bounding boxes by matching its Intersection Over Union (IOU) with selected regions from previous frames. Since the approach can consume frames online, early action recognition is naturally decided by the classification result from contemporary observations.

Arguably, human active regions (e.g., [80, 141]) are still very coarse if being compared with human skeleton analysis [169] where each body part (e.g., arms, elbows, shoulders etc.) is spatially spotted and used. Correspondingly, other work [14] has suggested to first use human body key-point extraction tools to locate an actor's body structure information and select the salient body feature parts for early action recognition, with a reinforcement policy learning [107] function for performing the selection. Along similar lines, research using 3D body skeleton analysis for early action recognition has focused on adaptively selecting the best historical information window for predictions in data-driven fashion [93].

Notably, the above reviewed works are mostly actor centric, i.e., they base their inference on analysis of the actor alone, with little regard for other objects and scene context. In contrast, a recent effort has instead considered object centric human actions (e.g., fetching objects or moving something to somewhere) [184]. This work reckoned that the key to reasoning about complex activities is the pattern of temporal developments between subactions and/or objects at various scales, something that was not modeled explicitly in previous work. In particular, the authors built a ConvNet to perform feature extractions and 2-layer MLPs to capture frame-wise temporal relations. More specifically, multiple sets of frames sampled from various temporal scales are combined together as the final representation to classify. For example, the collection of frame index groups $(1, 9)$, $(2, 5, 9, 12)$ and $(1, 4, 10)$ represent three sets of temporal relations where states of object or actor (i.e., human hand in Figure 2.3) transit differently in both speeds and conditions.

While the above approach reasons about temporal relations, it aggregates frame-wise information
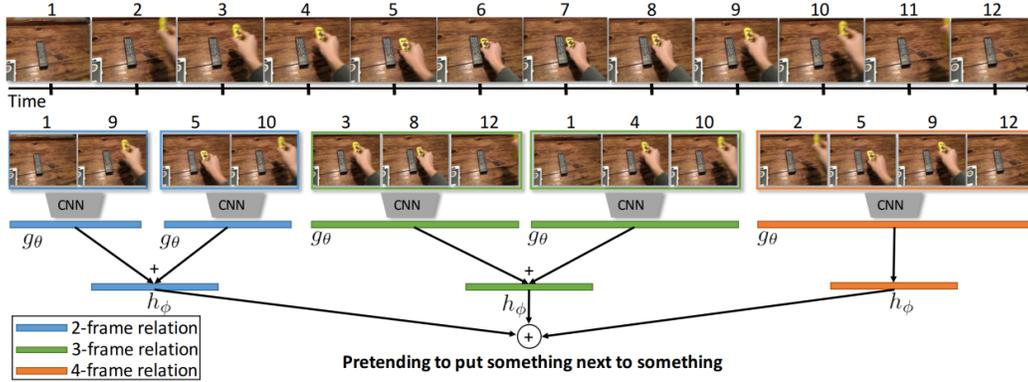
Figure 2.3: Example of temporal relation reasoning through sampling frames with multi-scale time intervals [184]. Frame group $(1, 9)$ denotes the relation **put next to** whereas group $(1, 4, 10)$ denotes the **take-away** relation. Considering all sub-relations leads to the inference of the action label. Figure reproduced with permission from [184].

globally and thereby does not exploit spatial relations of elements within a single frame. More recent work has extended this principle to the spatial dimension [149]. Similar to [141], it used a learning based detector to generate actor proposals for all characters in one frame. The relation between multiple actors is modeled as nodes and edges as in a graph [136] and a Graph Neural Network (GNN) [43] is adopted to diffuse information among nodes. Following this direction, a multi-scale spatio-temporal graph based method further explores fine-grained human-to-human and human-to-object interactions for this task [167].

Finally, some other work considers early action prediction as a special case of action recognition in that only a few action clips that are easily differentiable under full observations become difficult when incomplete. To deal with such difficult cases, they equipped the regular recognition framework with an extra memory module to log hard examples explicitly [66].

2.3 Knowledge distillation Based

Methods in the previous section mainly focus on reasoning about the future directly from the past. The future unobserved information provides negligible contribution to the partially observed information, except the overall action label being used for training. Here, it is worth remembering that previous work has shown that full observation almost always excels over the partial and no one has forbidden the use of unobserved information in training [129]. The more urgent question concerns *how* to use such unobserved information. Hereby, we discuss research that attempts to lend power from unobserved data in training in order to either enrich the feature representation of partial data or encourage the classifiers to recognize

partial data with ease.
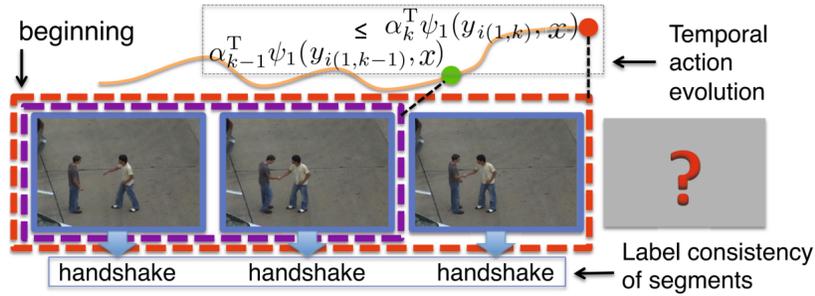
Figure 2.4: Temporal action evolution patterns can be enforced through a score function [68], as given in Equation 2.1. Classifiers may work better with partial videos by considering its ranking score with full videos. Figure modified with permission from [68].

Initial work that informally considered mutual information in early action recognition trained video segment classifiers that align the temporal order of partial observations to training data [64, 68]. They achieved this result by using a Structured SVM (SSVM) [155]. More specifically, they evenly divided action videos into equal-length segments, extracted features from each segment, integrated segments into one chunk, such as $y_{1:k}$, and trained SSVM classifiers according to the following constraints (without loss of generality, two progress levels $(k-1, k)$ is given):

$$\alpha_{k-1}^T \psi_1(y_{1:k-1}, x) \leqslant \alpha_k^T \psi_1(y_{1:k}, x), \tag{2.1}$$

where $y_{1:k}$ is the video segment feature accumulated from the very begining till the k_{th} level, x is the groundtruth label, α_k^T is the score matrix for progress level k and $\psi_1(y_{1:k}, x)$ is a joint feature map that represents the spatio-temporal features of action label, x, given a partial video $y_{1:k}$. The authors assumed 10 progress levels, such that $k \in (1, 10)$.

The above formulation emphasized temporal label consistency: Besides assigning a correct class label to given input videos, classifiers also need to capture the temporal evolution of actions by monotonically increasing the score for segments across time (first seen in early event detection [48]). The intuition coincides with daily human experiences: As we see more, we should know better than before. Some other work follows a similar intuition and built their early activity detection model based on ranking loss [99]. It also has been proposed to incorporate global human intentions as a complimentary learning objective to assist partial observation learning, but for action prediction in first-person-view videos [132]. One particular work further argued for assigning various weights for features from every progress level and suggested using temporal saliency to locate the corresponding weights [79]. Also, other work adopted the same monotonicity constraints and grounded the feature aggregation on human poses [120]. Finally, an

effort focused on refining the classification confidence score at each time-step in an incremental manner [49]. That work relied on an attention mechanism to provoke high confidence on early observations. Notably, all these efforts differ from methods in the previous section in that unobserved data gets involved to empower the partial data during training.

Yet another approach to using a scoring function was proposed by the originators of this direction for early action recognition [68]. In this alternative approach the attempt was made to infer complete future features from partial observations [69], something not previously attempted in early action recognition. Their work supposes that we have features of partial observation, $g(y_{1:t})$, and full observation, $g(y_{1:T})$ ($t \in (1, \ldots, T)$), in the training stage, and proposed to reconstruct the full from partial according to:

$$\underset{W,\theta}{\operatorname{argmin}} ||g(y_{1:T}, \theta) - W g(y_{1:t}, \theta)||^2, \tag{2.2}$$

where W is the feature transformation matrix and $g(y, \theta)$ is the feature extraction function with parameter θ. As shown in Figure 2.5, raw partial observations (e.g., riding on trails with green grass background) could be mixed with any other actions that share the similar background context (e.g., golfing or swimming in pools aside the lawn). Transforming the partial into full observations at the feature level helps differentiate between the possibilities (i.e., capturing more details of the bicycle structure as well as human riding motions).
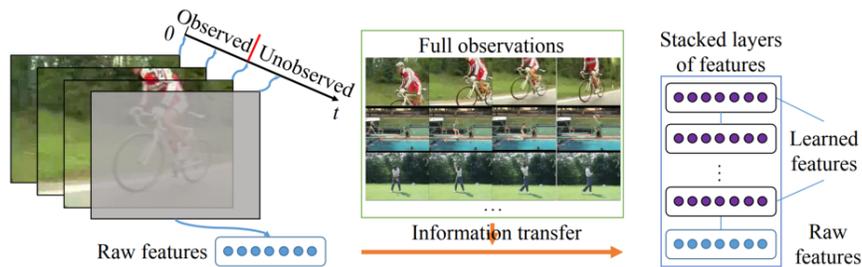
Figure 2.5: Illustration of the feature transformation approach of [69]. Raw incomplete features are transformed into features from the fully-observed to gain more discriminative information. Figure reproduced with permission from [69].

In implementation, the state-of-the-art 3D CNN video features (i.e., C3D [154]) were employed for g and the transformation matrix, W, was learned. A downside of this approach is that multiple transformation matrices needed to be learned, i.e., one for each k. A more recent upgrade on this work [70] overcame this concern by converting the deterministic feature transformation process [69] to a conditional generative process, with an adversarial learning component [35] to promote realism. The progress level, k, then plays as a hyper-parameter to inform the regressor (emboddied as a ConvNet) about the input

type and correspondingly adjust its weights for the desired output. A subsequent development argued for alleviating the heavy burden of high-dimensional feature regression [118]. The authors suggested transforming features into a low dimension binary code, as seen in Hashing encoding theory [9], and then operating the feature reconstruction on the binary encoded counterpart.

Other work approached the same problem by encouraging classifiers to assign correct class labels as early as possible [133]. They noticed that early work also made efforts on regularizing the classifiers [68]; however, none had explicitly enforced the early recognition of actions. In their framework, a sequence of video frames $(y_0, ..., y_{T-1}, y_T)$ (denoted as Fr_t in Figure 2.6) is processed by a ConvNet extractor, LSTM sequential aggregation and MLPs as the final classifier, which would regress a sequence of class probability for certain class index c as: $(\hat{x}_0(c), \hat{x}_1(c), \ldots, \hat{x}_T(c))$ (denoted as $y^t(k)$ in Figure 2.6).
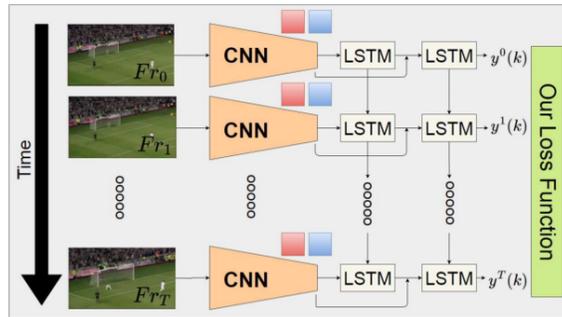
Figure 2.6: Algorithmic diagram of [133]. Sequential frames are processed by deep neural networks step by step and a time-variant cross-entropy loss function encourages early recognition. Figure reproduced with permission from [133].

They designed a novel cross-entropy loss to scale the false-positive loss value based on the incoming time index according to

$$\mathcal{L}(x, \hat{x}) = -\frac{1}{C} \sum_{c=1}^{C} \sum_{t=1}^{T} [x^t(c)log(\hat{x}^t(c)+ \frac{t(1 - x^t(c))}{T}log(1 - \hat{x}^t(c))],$$

(2.3)

where C denotes the total number of action classes.

Recalling that the false negative term (e.g., first term in Equation 2.3) aims for correctly assigning labels, whereas the false positive term (second term) for suppressing wrong labels, the authors made the relative weights for the former higher than the latter at the beginning of the sequence. That means high probability score for incorrect class labels is allowed as long as the score for correct ones is high enough. The key intuition behind this choice is tolerance for ambiguity at early stages. Therefore, the overall loss encourages the correct recognition in early time stamps.

Parallel research also exploited the property of loss functions for early action recognition [52]. Instead

of generating hard labels (e.g., SVMs only output one fixed label and discard other possibilities), they proposed to produce soft labels in a regression fashion, rather than classification. As a concrete example, they map input videos $y_{1:t}$ using transformation matrix W to a vector x that contains scores for all possible classes C as:

$$x = s(t)Wg(y_{1:t}, \theta), x \in \mathcal{R}^C, \tag{2.4}$$

where $g(y_t, \theta)$ is the feature extraction function, output x has dimension C and $s(t)$ is a scaling factor that monotonically increases along the time index. The authors adopted a real value time index, t, rather than the evenly allocated progress level, k, to enable online inference. Arguably, extensively adopted softmax classifiers have enabled probability scores for all classes and other related work has revealed its effectiveness [133]. However, no direct comparison of the performance between soft-regression and softmax classification was provided.
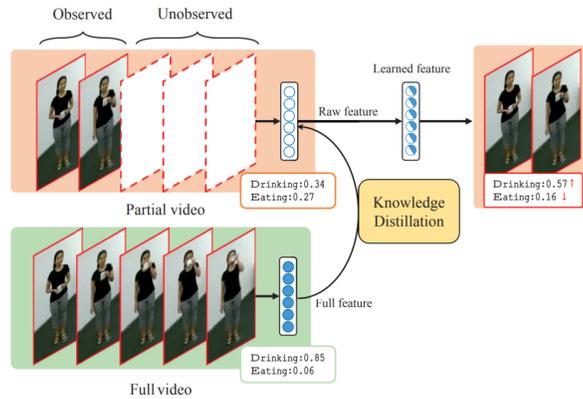
Figure 2.7: Illustration of the teacher-student learning approach. The student branch takes partial action clips as input (e.g., top row), while the teacher branch (bottom row) takes the full clip. Knowledge distillation implemented through feature distance minimization encourages the student branch to learn from the teacher. Figure reproduced with permission from [165].

Very recently, knowledge distillation [47] has been applied to squeeze fully-observed information into partial observations [165]. Two neural networks were employed: One acts as a teacher that gets access to full observations; the second acts as a student with only partial access. The student's representation is driven toward that of the teacher by minimizing the Mean Square Error (MSE) between the latent features and the Maximum Mean Discrepancy (MMD) between the predicted classification probabilities; see Figure 2.7. Some follow on efforts emphasized both the feature reconstruction consistency and semantic knowledge consistency via integrating multiple loss functions (e.g., feature distance loss, adversarial loss and classification loss) for a joint optimization [162,170]. Indeed, it seems that the distance metric plays a critical role in building the connection between video parts. Further, one recent work has found significant improvements via adopting the **Jaccard** similarity metric to learn the transition between partial and full

observations [25].

2.4 Propagation Based

Another way to exploit future information is by propagating the partial observation into the future. Typically, this type of approach extrapolates the visual representations into the future and then applies an existing classifer. The key difference between propagation and mutual information is whether or not the system performs **temporal extrapolations**. Mutual information methods from the previous section emphasize how to enrich partial observations with future information, whereas propagation based methods in this section focus on how to generate an accurate future based on current observations. Mathematically, propagation based approaches start from any incomplete input $y_{1:t}$, assuming $t < T$, and produce the subsequent information $(y_{t+1}, y_{t+2}, ..., y_T)$; thus, they can be considered as temporal extrapolation. In contrast, during either the testing or training phase, mutual information based approaches do not extrapolate. Classification is then based on the the concatenation of the observed and extrapolated.

An early exemplar of the propagation approach operated by working with relatively high-level visual representations so that they can be readily processed by extant recognition tools [160]. In practice, they
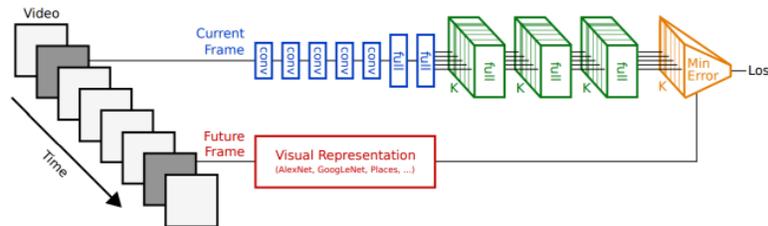
Figure 2.8: Deep regression network from [160] with mixtures of output layers. Blue boxes are convolutional and fully connected layers for unified processing on each frame, while green boxes represent the mixture of output layers. To push the predicted visual representation close to the real future, MSE loss (i.e., yellow box) is calculated with the groundtruth extracted from future frames in a self-supervised way. Figure reproduced with permission from [160].

formulated a deep regression network, consisting of multiple convolutional and fully connected layers, to temporally regress the future high-level deep feature (i.e., fc7 activation output from AlexNet [74]). To generate the label for future representations without supervision, they collected a small group of labeled video representations, calculated the similarity score and chose the best one. In order to model multi-modality, they extended the original single output structure to K mixtures (e.g., denoted as green boxes in Figure 2.8).

This approach has two positive attributes. First, to the extent that the deep features capture se-

mantically meaningful representations, they should be well suited to support recognition. Second, the mixture output helps to model the possibly multimodality of future outcomes. A notable downside is the potentially prohibitive dimensionality of the representation. Nevertheless, it unveiled a novel approach and inspired several follow-on efforts. A subsequent work used Recurrent Neural Networks (RNNs) with learnable Radial Basis Function (RBF) kernels to improve the nonlinearity of predictions [139]. Yet other work upgraded the deep regression component with a Markov Decision Process (MDP) as well as an RNNs [179], but shared the same intuition.

While the initial studies on propagation-based early action recognition achieved promising results, critical questions were left unanswered: What type of feature works best for propagation? How can video redundancy be dealt with? How can remote predictions avoid degeneracy? Further, design decisions appeared to be made heuristically without support from ablation studies. These concerns were addressed in an approach that made use of classic recursive filtering coupled with more careful design [182]; see Figure 2.9.
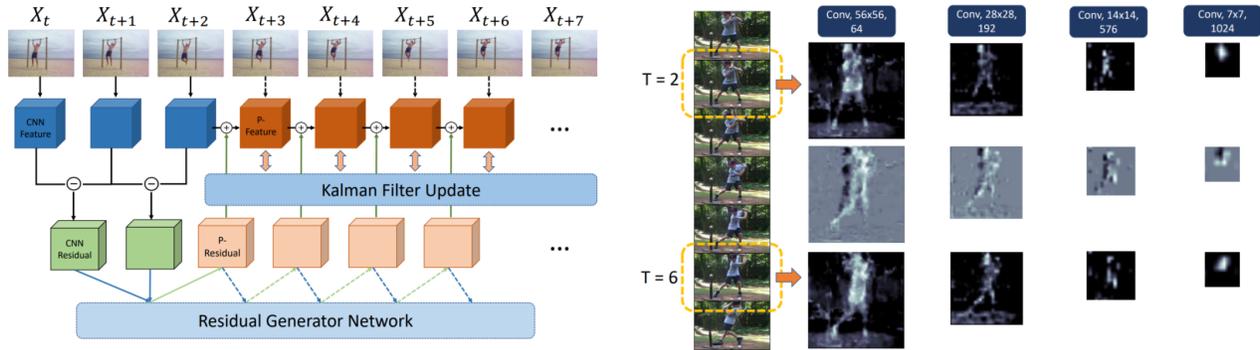
Figure 2.9: The left sub-figure depicts the overall workflow from [182] that revisited Kalman filtering under deep network structure: Partial observations (i.e., blue boxes) and their pair-wise residuals (i.e., green boxes) are used to propagate information into the future. Kalman filtering update is performed at each time step to combat error accumulation; The right sub-figure shows an ablation study that examines different feature stages through visualization. The intermediate feature stage (e.g., 3^{rd} column from left) performs the best since it keeps critical information of actors while suppresses noisy background present in the shallow layer. In comparison the vectorized features deeper in the network are too coarse and suppress important details. Numerical studies backed-up the visualizations. Figure reproduced with permission from [182].

They proposed to propagate the frame-wise feature residuals to lessen the temporal redundancy and thereby forced the model to focus on the underlying dynamics, rather than static background. Moreover, to combat error accumulation, Kalman filtering is incorporated into a deep network, which iterated the **prediction** and **update** steps to stablize long-term predictions. Further, they provided an answer to the above question regarding proper features by examining various feature layers in a deep network in

propagation and determined that the intermediate layer features, rather than vectorized last layer feature typically used in earlier work, is more suitable for propagation. Presumably, this intermediate level of representation provides the right level of abstraction from the raw video without loosing too much detail.

Other work combined the idea of feature propagation with a reinforcement reward function to encourage early recognition [32]; Yet other work propagated both the RGB and motion (i.e., optical flow) features from deep networks and adopted adversarial learning to boost the realism of their outputs [31]. In contrast to most recent work that heavily used deep features, an effort has employed hand-crafted Dynamic Image (DI) [7] features for propagation, but still depends on deep recognition models for the label assignment [124]. Finally, an approach that propagated both graph node features (e.g., 2D coordinates) and structures (e.g., adjacency matrices) has been proposed to tackle a recently appeared task, namely group-player early action recognition [13].

2.5 Datasets and Performance

In this section, we describe datasets, evaluation metrics and performance for the majority of early action recognition approaches discussed in this section. As all datasets are used for both action recognition and early recognition, comprehensive comparisons of them can be found elsewhere [59, 65].

Evaluation on early action recognition is based on the same recognition accuracy metric as in action recognition, but with a focus on the **observation ratio** K ranging from 10% to 100%, as discussed in the beginning of this chapter. Since each video only owns a single label, the video level mean average precision is reported against each K value.

UCF-101 [145] is a dataset collected from a real-life video platform (Youtube) and trimmed for action recognition (each video contains exactly one action). It includes 101 distinct action classes and 13,320 overall video clips with at least 100 videos for each category. UCF-11 and UCF-50 are two earlier versions of the same dataset that contain 11 and 50 action categories respectively. One critical issue of UCF-101 is that videos from YouTube could be very biased by low-level features, meaning low-level features (i.e., color and gist) are more discriminative than mid-level features (i.e., motion and shape). Researchers have been progressively evaluating both action recognition and early action recognition on UCF-101.

HMDB-51 [77] is a large scale human action recognition dataset that comprises 51 daily action categories. It is worth noticing that it contains some fine-grained human facial motions, such as smiling, laughing, chewing etc, in static background windows, which are not seen in other comparable datasets. This dataset also encompasses a wider range of camera motions and viewpoint changes compared to UCF-101. In total, there are a total of 6,766 video clips with at least 102 videos for each class. There are

three official data splits and the averaged evaluation score is often reported as final results.

JHMDB-21 [57] is a subset of HMDB-51 with a special focus on 21 human body joint action categories where additional annotations on joints (13 2D-coordinates body joints) are provided. Compared with HMDB-51, categories that mainly contain facial expressions (e.g. smiling), interaction with others (e.g. shaking hands), and very specific actions (e.g. cartwheels) were excluded. Therefore, there are a total of 928 video clips.

BIT [67] consists of 8 classes of human interactions (bow, boxing, handshake, high-five, hug, kick, pat, and push), with 50 videos per class. Videos are captured in realistic scenes with cluttered backgrounds, partially occluded body parts, moving objects, and variations in subject appearance, scale, illumination condition and viewpoint. Even though BIT has a relatively limited number of classes and videos, it is a complex dataset in that: **1.** The backgrounds as well as actor appearances are highly non-discriminative. Various categories often share the same background and actors; **2.** The starting and ending phases of the videos are highly similar (i.e., actors standing still and facing each other). Thus, researchers often have to resort to motion based features for comparable performance.

UT-Interaction (UTI) [131] is comprised of 2 sets with different environments. Each set consists of 6 types of human interactions: *handshake, hug, kick, point, punch and push.* Each type of interaction contains 10 videos, to provide 60 videos in total. Videos are captured at different scales and illumination conditions. Moreover, some irrelevant pedestrians are present in the videos.

TV Human Interaction [115] consists of 300 video clips collected from over 20 different TV shows. It contains five action classes: *handshake, high five, hug, kiss and none.* The class "none" represents all other more general actions such as walking and standing. Annotations are provided for every frame of the videos, including the upper body bounding boxes, discrete head orientations and action labels for each person.

Something2something [36] is a dataset that shows human interaction with everyday objects. In the dataset, humans perform a pre-defined action with a daily object. It contains 108, 499 video clips across 174 classes. The dataset enables the learning of visual representations for physical properties of the objects and the world. Compared with others, the dataset shows primarily human related motions (mostly only hand motions are visible).

SYSU-3DHOI [51] (3D Human-Object Interaction) includes 480 RGBD sequences from 12 action categories, including "playing phone", "calling phone", "pouring", "drinking", etc. For building this set, 40 participants were asked to perform 12 different activities freely. For each activity, each participant manipulates one of the six different objects: phone, chair, bag, wallet, mop and besom. Therefore, there are in total 480 video clips collected in this set. The contained activity samples have different durations,

ranging from 1.9s to 21s. For each video clip, the corresponding RGB frames, depth sequence and skeleton data were captured by a Kinect camera [181].

NTU-RGBD [137] is by far the largest public set for 3D action recognition and prediction. It contains more than 56,000 video samples with about 4 million frames from 60 action categories. All of these action samples were recorded by a Kinect camera [181] from three different views. For collecting this set, 40 subjects were asked to perform certain actions several times. This set is very challenging for early action prediction mainly due to its larger scales of quantity, greater diversity in action categories and more complexity in human-human interaction and human-object interaction.

ORGBD [176] (Online RGB-D) was collected for online action recognition and early action prediction. There are seven types of actions that people often do in the living room: drinking, eating, using laptop, picking up phone, reading phone (sending SMS), reading book, and using remote. All these actions are human-object interactions. The bounding box of the object in each frame is manually labelled.

Sports-1M [60] contains 1, 133, 158 video URLs, which have been annotated automatically with 487 labels. It is one of the largest video datasets. Very diverse sports videos are included in this dataset, such as shaolin kung fu, wing chun, etc. The dataset is extremely challenging due to very large appearance and pose variations, significant camera motion, noisy background motion, etc. However, since the videos originate from Youtube and authors have not maintained the complete video set, some URLs become invalid or expired. Moreover, the automated labelling is known to have yielded incorrect groundtruth in some cases.

Jester [103] is a recent video dataset for hand gesture recognition. It is a large collection of densely-labeled video clips that shows humans performing pre-defined hand gestures in front of a laptop camera or webcam with a frame rate of 30 fps. There are in total 148,092 gesture videos under 27 classes performed by a large number of crowd workers. The dataset is divided into three subsets: training set (118,562 videos), validation set (14,787 videos), and test set (14,743 videos).

OAD [89] was captured with a Kinect camera [181] in daily-life indoor environments and thus provides RGBD data. It includes 10 actions. The long video sequences in this dataset correspond to about 700 action instances. The starting and ending frames of each action are annotated in this dataset. 30 long sequences are used for training, and 20 long sequences are for testing.

PKU-MMD [92] is a large dataset for 3D activity analysis in continuous sequences. The data set contains 1076 long video sequences in 51 action categories, performed by 66 subjects in three camera views. It contains almost 20,000 action instances and 5.4 million frames in total. This dataset also provides multi-modality data sources, including RGB, depth, Infrared Radiation and Skeleton.

OoPS! [20] is proposed recently to study unintentional human actions. It includes a variety of daily human activities. Instead of focusing on action classification, the paper pays special attention to whether the performer succeeds/fails at achieving the goal. The dataset contains 20,338 videos and only has three class labels: Intentional, Transition and Unintentional. This paper conducted the action-intention prediction task on this dataset: early predicting the failure before it happens (1.5 seconds into future).

	Year	No. Videos	No. Actions	Avg. Len	Domain	Data-Modality
OoPS! [20]	2020	20,338	3	**9.4s**	**Action-intention**	RGB, Pose, Flow, Naaration
Jester [103]	2019	146,092	27	**3s**	**Gesture**	RGB
Something2something [36]	2017	108,499	174	**4.03**	**Activity**	RGB
PKU-MMD [92]	2017	1076	51	-	**Activity, HOI**	RGB+D, IR, Skeleton
OAD [89]	2016	59	10	**219**	**Activity**	RGB+D, Skeleton
NTU-RGBD [137]	2016	56,000	60	-	**Activity, Interaction**	RGB+D, IR, 3D-Skeleton
SYSU-3DHOI [51]	2015	480	12	-	**Activity, HOI**	RGB+D
UT-Interaction [131]	2015	60	6	**3s**	**Human Interaction**	RGB
Sports-1M [60]	2014	1,133,158	487	-	**Action**	RGB
ORGBD [176]	2014	-	7	**12**	**Action, HOI**	RGB+D, Skeleton, BBx
JHMDB-21 [57]	2013	928	21	-	**Action**	RGB, Flow, Skeleton, Contour
UCF-101 [145]	2012	13,320	101	**6.39s**	**Action**	RGB
BIT [67]	2012	50	8	-	**Human Interaction**	RGB
HMDB-51 [77]	2011	6.766	51	-	**Action**	RGB
TV Human-Interaction [115]	2010	300	20	-	**Human Interaction**	RGB

Table 2.1: Summary of datasets for early action recognition.

Performance evaluation of early action recognition has been focused on consideration of four of the above datasets: UCF-101, J-HMDB, UT-Interaction and BIT. We restrict performance comparisons to that set. For J-HMDB as shown in Table 2.2, we follow the standard procedure to report the recognition accuracy based on the first 20% frames of videos. Some work also reports results regarding the whole set of observation ratios and we refer readers to [141, 170] for more details. As noticed, recent advances based on motion extrapolation and Kalman filtering [182] as well as an RNN with RBF kernels [139] have pushed the early recognition accuracy using only 20% frames to around 75%. Thus, there is still room for further improvements. Very recently, an effort achieved significant improvements (i.e., 83%) via feature mapping learned through the Jaccard distance [25]. On the other hand, for the UT-Interaction dataset, recent approaches can achieve over 90% accuracy with only the first 20% frames and can obtain almost perfect accuracy over the first 50%; see Table 2.3. For this dataset, a wide variety of approaches perform very well. These observations somewhat indicate that UT-Interaction is saturated. Similarly, the UCF-101 dataset has been demonstrated to be easily recognizable using less than 30% frames, as shown in Table 2.4. Again, a wide variety of aproaches are able to do well on this dataset. In contrast, the BIT dataset, for most approaches, still provides difficulties for early recognition, probably due to its limited size. Although, the most recent approaches can do very well having viewed 50% or more of the data.

Table 2.2: Summary of J-HMDB Performance

Methods	Accuracy @ 20%
Within-class Loss [99]	33%
DP-SVM [144]	5%
S-SVM [144]	5%
Where/What [143]	10%
Context-fusion [56]	28%
ELSTM [133]	55%
G. Singh's [141]	59%
TPnet [140]	60%
DR2N [149]	66%
Pred-GAN [170]	67%
RBF-RNN [139]	73%
RGN-KF [182]	78%
JVS+JCC+JFIP [25]	83.5%

Table 2.3: Summary of UTI Performance. Notice that some work reports results for set 1 and set 2 of the UTI dataset separately e.g., [14] , while others report the averaged results.

Methods	Accuracy @ 20%	Accuracy @ 50%
S-SVM [144]	11.00	13.40
DP-SVM [144]	13.00	14.60
CuboidBayes [129]	25.00	71.00
CuboidSVM [130]	31.70	85.00
Context-fusion [56]	45.00	65.00
Within-class Loss [99]	48.00	60.00
IBoW [129]	65.00	81.70
DBoW [129]	70.00	85.00
BP-SVM [82]	65.00	83.30
ELSTM [133]	84.00	90.00
Poselet [120]	-	73.33
PA-DRL (set 1) [14]	69.00	91.70
PA-DRL (set 2) [14]	58.00	83.30
Future-dynamic Image	89.20	91.90
Early-GAN [162]	38.00	78.00
Pred-GAN (set 1) [170]	-	100.00
Pred-GAN (set 2) [170]	-	85.70
RBF-RNN [139]	-	97.00
Joint-GAN [31]	98.30	99.20
SPR-Net [49]	-	85.30

Table 2.4: Summary of UCF-101 Performance. Notice that some work only reports accuracy at a few ratios (i.e., 50% and 100%). The symbol ≈ means approximated number from the original plotted graph, rather than actual number, due to the unavailability of some results.

Metrics	Obervation Ratio									
Methods	10%	20%	30%	40%	50%	60%	70%	80%	90%	100%
IBoW [129]	36.29	65.69	71.69	74.25	74.39	75.23	75.36	75.57	75.79	75.79
DBoW [129]	36.29	51.57	52.71	53.13	53.16	53.24	53.24	53.24	53.45	53.53
MTSSVM [68]	40.05	72.83	80.02	82.18	82.39	83.21	83.37	83.51	83.69	82.82
MSSC [10]	34.05	53.31	58.55	57.94	61.79	60.86	63.17	63.64	61.63	61.63
DeepSCN [69]	45.02	77.64	82.95	85.36	85.75	86.70	87.10	87.42	87.50	87.63
Mem-LSTM [66]	51.02	80.97	85.73	87.76	88.37	88.58	89.09	89.38	89.67	90.49
MSRNN [80]	68.00	87.39	88.16	88.79	89.24	89.67	89.85	90.28	90.43	90.70
PA-DRL [14]	-	-	-	-	87.30	-	-	-	-	87.70
AAPNet [70]	59.85	80.85	86.78	86.47	86.94	88.34	88.34	89.85	90.85	91.99
Student-Teacher [165]	83.32	87.13	88.92	90.85	91.04	91.28	91.28	91.23	91.31	91.47
Joint-GAN [31]	-	84.20	-	-	85.60	-	-	-	-	-
RGN-KF [182]	83.12	85.16	88.44	90.78	91.42	92.03	92.00	93.19	93.13	93.13
Pred-CGAN [170]	-	-	-	-	91.14	-	-	-	-	93.76
Eearly-CGAN ≈ [162]	78.00	84.50	86.00	88.00	89.40	89.95	90.01	90.53	90.76	91.05
SPR-Net [49]	88.70	-	-	-	91.60	-	-	-	-	91.40
STGCN [25]	80.26	-	89.86	-	92.87	-	94.08	-	94.43	-
JVS+JCC+JFIP [25]	-	91.70	-	-	-	-	-	-	-	-

Table 2.5: Summary of BIT-101 Performance

Metrics	Obervation Ratio									
Methods	10%	20%	30%	40%	50%	60%	70%	80%	90%	100%
IBoW [129]	22.66	24.22	37.50	48.44	48.44	52.34	46.09	49.22	42.97	43.75
DBoW [129]	22.66	25.78	40.63	43.75	46.88	54.69	55.47	54.69	55.47	53.13
MTSSVM [68]	28.12	32.81	45.31	55.47	60.00	61.72	67.19	70.31	71.09	76.56
MSSC [10]	21.09	25.00	41.41	43.75	48.44	57.03	60.16	62.50	66.40	67.97
DeepSCN [69]	37.50	44.53	59.38	71.88	78.13	85.16	86.72	87.50	88.28	90.63
MSDA [15]	-	-	-	-	70.00	-	-	-	-	81.50
GLTSD [79]	26.60	-	-	-	79.40	-	-	-	-	-
PA-DRL [14]	-	-	-	-	85.90	-	-	-	-	91.40
AAPNet [70]	38.84	45.31	64.84	73.40	80.47	88.28	88.28	89.06	89.84	91.40
RGN-KF [182]	35.16	46.09	67.97	75.78	82.03	88.28	92.19	92.28	92.16	92.16
Pred-CGAN [170]	-	-	-	-	87.31	-	-	-	-	91.79
Eearly-CGAN ≈ [162]	38.80	58.00	79.00	90.00	96.00	96.00	96.00	96.00	96.00	96.00
SPR-Net [49]	84.10	-	-	-	100.00	-	-	-	-	99.20
STGCN [25]	46.09	-	58.59	-	81.25	-	89.06	-	86.72	-

Chapter 3

Future Action Prediction

3.1 Overview
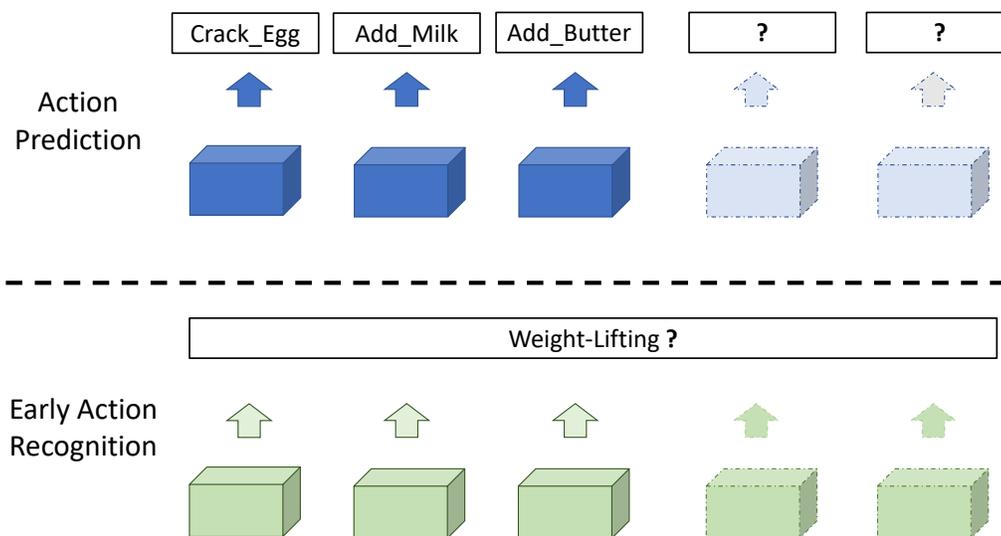
Figure 3.1: Illustration of the general dissimilarity between future action prediction and early action recognition. The latter (bottom row) aims to infer a class label for an ongoing action from partial initial observations; the former aims to anticipate (strings of) future actions from initial observations.

Other than recognizing certain *single* actions at its early stage, another emerging topic has caught researchers' attention, that is future action prediction. It differs from the previously discussed problem, early action recognition of Section 2.1, in that its goal is to produce predictions of subsequent future actions. As illustrated in Figure 3.1, the input (as solid boxes) and its output (as boundary-dotted boxes) in Early Recognition (bottom row) share the same label, such as **Weight-Lifting**, whereas in Action

Prediction (top row) input and output are often different yet logically related.

The ability to deduce possible future events from previous and current observations is one of the core functionalities needed by contemporary intelligent systems. In the example of automated traffic systems, i.e., artifical integellience guided vehicles, every driver needs to be cautiously aware of surroundings and would benefit from smart prediction systems that can tell the specific next action of in-coming vehicles (e.g., turning-right), to better plan its own motion in response (e.g., slow-down or keep driving). Early action recognition does not satisfy such a demand, as it simply predicts the completion of an in progress action and thereby lacks sufficient prediction horizon to anticipate future actions. Exploring the plausible future is well studied in other fields, e.g., weather forecasting and stock price prediction. In the computer vision community, however, researchers have only recently become heavily active in exploring solutions.

To examine future action prediction in detail, we group existing research into three major divisions: **1.** activity singleton or sequence prediction; **2.** joint prediction of activity semantics and times; **3.** ego-centric oriented action prediction. The first division discusses the most frequent setting for future action prediction: given the raw video frame observations, $y_{1:t}$, and/or its action semantics, $x_{1:t}$, produce the semantic label for actions that would happen afterward. The prediction can be either the immediate single following action, i.e., x_{t+1}, or a sequence of future actions, i.e., $x_{t+1:t+n}$. Note that the prediction provides only action semantics (e.g., labels); time of occurrence and duration are left ambiguous. The second division considers not just prediction of sequences of future actions, but also their times. The last division tries to assist action forecastings in ego-centric settings, where observed frames are taken from first-person-view and thus mimic real life carry-on camera settings.

3.2   Activity Singleton or Sequence Prediction

This section describes the general version of future action prediction, which is mapping the video input, that often contains a sequence of observed actions, into its subsequent sequence of actions, as shown in Figure 3.1 top row. The length of the predicted sequence can vary from one to many.

Hereby, we formally define the problem according to one of the initial works in future action prediction [12]. Let a sequence of observed video frames be given as $y^{obs} = (y_1, y_2, \ldots, y_t)$ and its corresponding frame-wise action labels be given as $\hat{x}^{obs} = (x_1, x_2, \ldots, x_t)$. Note that the symbol x^{obs} represents the groundtruth video level labels annotated through human labor. However, in most studies, researchers would estimate the action labels with existing methods, rather than directly using the groundtruth video level labels. To clarify this nuance, we denote the estimated action labels as $\hat{x}^{obs}$. Since actions can take large time spans, the frame-wise action labels can be either repeating or idle (no activity performed). Our

task is to produce the future action semantics $x^{unobs} = (x_{t+1}, x_{t+2}, ...)$ that happen at unobserved time indices, i.e., $t+n, n > 0$. The set of future time indices of interest, $N = (n, n+1, n+2, ...)$, are referred to as the prediction horizon. So in other words, the task is to seek a function f that maps $y^{obs} \to x^{unobs}$ with specified prediction horizon n_i.

Indeed, compared with the standard video action recognition definition $f : y^{obs} \to x^{obs}$, the above definition seems like a temporal shifted version of the recognition function, e.g., shifting x^{obs} to x^{unobs} by a desired time span. However, it is worth noticing that the semantic level information of observed video input, y^{obs}, can be inferred beforehand, either with groundtruth annotations or extraction algorithms, and therefore it is possible to operate a high-level semantic mapping $y^{obs} \to \hat{x}^{obs} \to x^{unobs}$ as a solution. Actually, most researchers adopted the second approach, especially for long-term predictions that typically cover multiple subsequent actions.

For example, early work along this line [12] proposed to treat any sequence of actions performed by all actors within a video as connected nodes in a Markov Random Field model (MRF) [33], as shown in Figure 3.2. These nodes are connected with a manually defined edge potential function that reflects a reasonable range of spatiotemporal affinity between any two nodes. Furthermore, a node potential function that denotes how likely a certain observation, y_t, belongs to certain action label, x_t, was implemented as SVM classifiers. More specifically, the observed action labels, $\hat{x}^{obs}$, come from a pre-trained recognition system
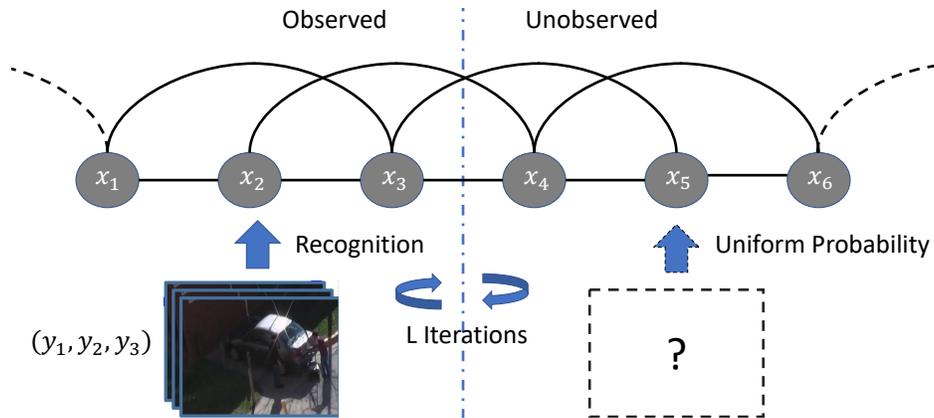
Figure 3.2: Workflow diagram of MRF based approach from [12]. Some nodes are initialized as recognized activity labels from observed video frames (e.g., x_1, x_2, x_3). Other nodes are initialized with uniform probability for entire action candidate set (e.g., x_4, x_5, x_6). Loopy belief propagation (LBP) will run L iterations to pass messages between nodes. As noted, the direct mapping from the past to the future is built upon high level action semantic labels $\hat{x}^{obs} \to x^{unobs}$, instead of raw frame inputs. The labels are extracted from raw frame inputs using a standard action recognition algorithm. Figure modified with permission from [12]

.

that relies on a Bag-of-Word (BoW) approach over Space Time Interest Points (STIP) and Multi-class

SVMs. To do the inference of the next action labels x_{t+n}, loopy belief propagation (also known as the sum-product algorithm) [109] is performed to compute the conditional distribution: $p(x^{unobs}|\hat{x}^{obs})$. The capacity of a Markov model greatly depends on the order of dependency. The adopted MRF model in [12] specifies the maximum size of the neighboring clique as four nodes (two nodes in the past and two nodes in the future, if any), which leads to faster inference but limited performance.

Another research effort also adopted this high-level label inference scheme but relies on a more flexible Markov model, namely the Various order Markov dependency Model (VMM) implemented as graph (or tree) structures to model the temporal relation among activity sequences, as shown in Figure 3.3 (a). Work that relies on such structures for future activity parsing include [42, 86, 87, 116]. These approaches advocate to represent activity sequence information as a sequence of ordered discrete alphabet symbols (grammars) and group them with n-gram representations. Again, to obtain action labels from observed inputs, off-the-self action detection as well as recognition algorithms are needed.

As an example, some efforts choose to build a Probabilistic Suffix Tree (PST) for representing complex activity structures [86]. The PST structure is learned from training data, with the pretrained tree providing the chained transition probabilty for each suffix. Thus, the inference of future actions can be viewed as finding the most probable suffix leaf node given prefix node inputs, as seen in Figure 3.3 right side. The obvious drawback of Markov based methods is that the length of prediction is pre-defined in

(a)   (b)
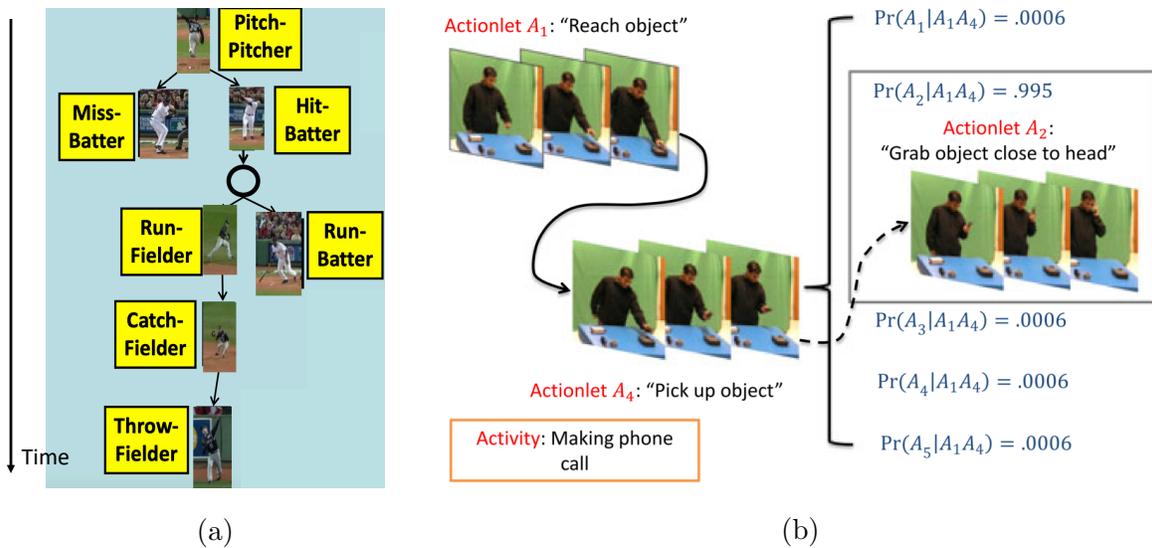
Figure 3.3: (a) Example of representing an activity sequence as a tree graph from [39]. The temporal development is structured as tree nodes and leafs in a top-down fashion, where the top stands for the starting and bottom for ending. (b) The prediction of future action "Grab object close to head" is selected by traversing the PST given prefix input and choosing the most likely suffix. Figure reproduced with permission from [86].

the training set. For example in [86], the height of the PST is decided during training and thus the model

can not be easily extended to undefined future lengths. However, the *n-gram* representation of video activity is an over-simplified method, as critical information contained in image space is ignored. To this end, a recent effort researched on video representation learning in Riemannian geometry space, where the learned model can enable hyperbolic space embeddings, which is a continuous version of tree structures as shown in Figure 3.3. Thus, the predictive model supports building activity hierarchies for both early and future action prediction [150].

Instead of using generative probabilistic models as above, a discriminative instance of Markov Chains, Conditional Random Fields (CRFs) [78], has been suggested to combine with particle filters for anticipating future human-object interactions [71]. This combination enables a straightforward temporal extrapolation, as the particle filters are a type of recursive Bayesian filter that can propagate into unlimited time and space. In particular, they jointly modeled the future activity semantics with object affordances as well as movement trajectories. Notably, a recent work has advanced generative adversarial grammar learning for activity forecasting [117].

The aforementioned approaches build models on a high-level semantic space, rather than raw video input or even visual features. The underlying assumption is that activity labels are linked by certain standard rules, e.g., actions for making certain salads are dictated by their instructions. Therefore, straightforward modelling on sequences of action labels is sufficient. Yet in real life, scenarios exist where future actions are less well predicted by well defined scripts. For example, pedestrians walking can happen in a wide range of places and times. It can lead to "road crossing", "turning", "stop and watch", etc. So, reasoning the future from an extremely common action is hard. Furthermore, in surveillance camera settings, actors can enter the field of view at times that obscure their previous actions, in which cases observing one's full activity sequence is impossible. To solve such problems, many researchers resort to a direct mapping from y^{obs} to future actions x^{unobs}.
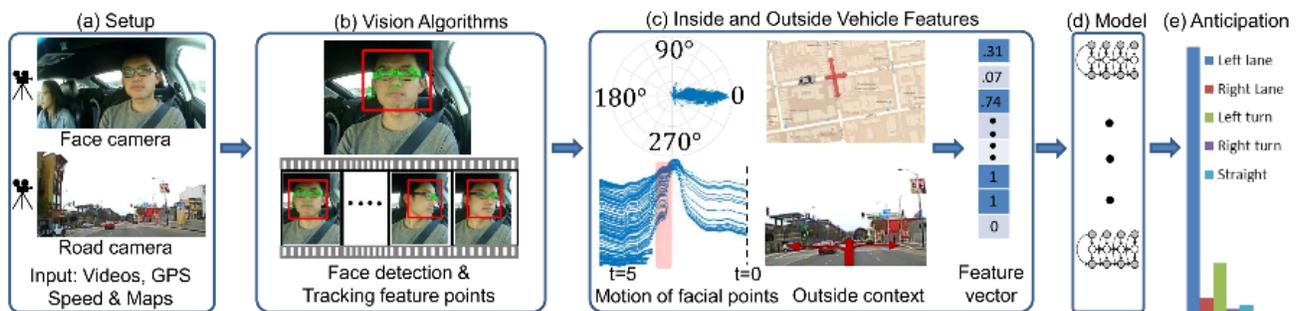

Figure 3.4: The setting of [55] combing multiple cues. Multiple cues collected from both inside and outside sensors are used to anticipate the next maneuver of drivers. Figure reproduced with permission from [55].

One early effort that followed this scheme proposed to rely on cues distilled from multiple visual sources

to predict the next driver maneuver [55]. Specifically, it collected information from four information sources: 1. facial motion trajectories and head angles detected from the driver-facing camera and feature tracking techniques; 2. speed logger from the vehicle system; 3. outside context from the road-facing camera; 4. the GPS map. They fused these four types of features with a novel HMM model where the maneuvers of drivers play as the hidden variables, outside sensory as control units and inside sensory as observations. However, since HMMs can not handle very high dimensional data, they converted most sensory cues into binary features, e.g., if vehicle speed is larger than 15km/hr, set the feature flag as 1 otherwise 0.

To encourage more expressive features, e.g., 3D human head/facial models, the authors made an upgrade on the sequential modelling module [56]. They replaced the HMMs with a recurrent neural network (an LSTM) and modulated four sensory channels with two individual RNNs. The final fusion was done by MLPs, followed by softmax classifiers. A critical difference between this work and earlier related approaches based on high-level intermediate inferences, e.g., [86], is the encompassed time span. Observations in the high-level based usually cover multiple actions that can takes **minutes**. In contrast, [56] only consumes frames that are 0.8 second (about 20 frames) ahead of the event happening, which has the potential to be deployed in real-time systems.

Similar work that focuses on short-term (or singleton) action prediction directly from past raw video inputs appears frequently in many sub-areas. One such effort proposed to leverage environment-agent interactions to anticipate risk actions as well as localize future risk regions [178]. Another worked on anticipating traffic accidents from vehicle cameras through stacking local detected object features and global context [151]. Some other work built a graph Convolutional Network (GCN) to model interactions between multiple agents and contextual objects to forecast future activities along with future trajectories of pedestrians [91]. Pedestrian road-crossing intention prediction is one concrete example of action prediction in traffic scenarios [121–123]. Again, in these works, past observations contain little or no evidently useful activity semantics, instead researchers found it helpful to mine cues from subtle details. As an example, an effort on anticipating crossing of pedestrians fused human poses, human visual representation extracted from deep networks, scene contextual information and vehicle speed, with a RNN, to make predictions [123].

3.3 Joint Prediction of Activity and Time

Joint prediction of future actions and their times has recently become an active research area for its importance in critical applications like automated navigation, human-computer interaction as well

as for the intellectual challenges it presents. All video-based action prediction work forecasts upcoming activities, here, the distinction is that both the activities and their times are predicted. An example is provided in Figure 3.5.
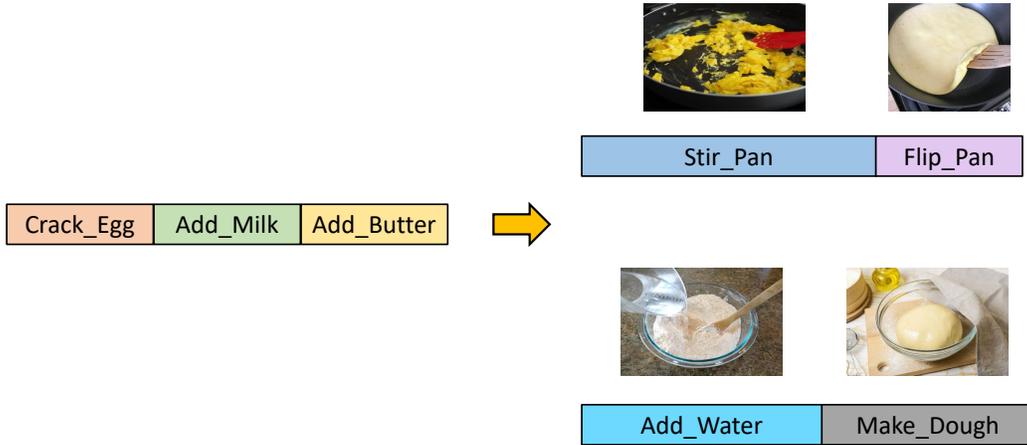
Figure 3.5: Predicting future activity labels and time durations. The left hand side shows initial observed activities, indicated as labelled boxes, with their temporal duration encoded by box length. The right hand side shows possible predictions.

An initial effort in this direction [100] adopted the Poisson process [62] as a key technique for activity inter-arrival time modelling. Notice that inter-arrival time differs from the absolute activity occurrence time in that it stands for the duration between the starting and ending of a certain action. The probability density function for the possible next action interval, τ, under the assumed Poisson distribution is

$$p(T = \tau) = \frac{\lambda^{\tau}}{\tau!}e^{\lambda} \tag{3.1}$$

where λ is the intensity function that controls the shape of the Poisson probability density function and its physical meaning is the average number of events among the whole time span. The Poisson distribution and sampling are plotted at Figure 3.6 (a) and (b) respectively.

There are three properties of the Poisson distribution that are worth noticing: **1.** Its samples are always positive integers. Equation 3.1 incurs a **factorial** function on time τ, thus its result must be an integer; additionally, by default researchers assume there is no negative time value, so that $\tau > 0$. To summarize, the Poisson distribution is a positive **discrete** function as shown in Figure 3.6 (b). **2.** Its shape is unimodal. **3.** It is memoryless in that the sampling of the current step is independent from any previous samples.

To support a stochastic modelling of time, the authors pre-defined a Gaussian Process prior for $\lambda(t)$ and learned its parameters (i.e., the mean and variance for certain actions and times) from training data

(a) (b)
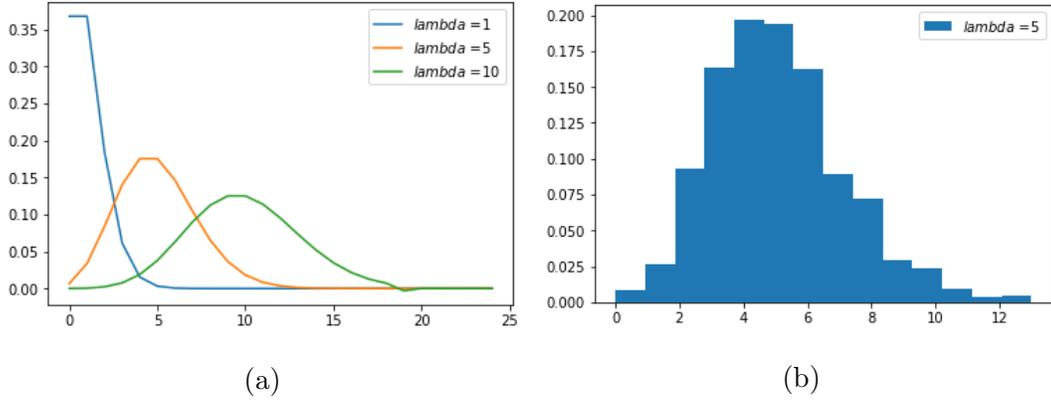
Figure 3.6: (a) Probability density function of Poisson distribution regarding various parameters λ. (b) Histogram distribution of 1000 sampled inter-arrival times from Poisson process with $\lambda = 5$. Since Poisson is a member of the exponential family, its distribution looks similar to a Gaussian distribution, but its data points are always positive integers.

for every action. Then, they performed testing through importance sampling. Unfortunately, they only train and test on a single action input and single action output, without considering recursive temporal predictions.

An extension of the initial work [100] used the temporal point process for the same task [105], where the intensity parameter λ is dynamically calculated from past observations as

$$p(T = \tau) = \lambda(\tau|x_{1:k})e^{-\int_0^{t|x_{1:k}} \lambda(t)dt}, \tag{3.2}$$

where $x_{1:k}$ represents past information, such as raw frames, semantic labels or past action times. Their model also supports stochastic time and action prediction via sampling the parameter λ through a Variational Auto-Encoder (VAE) embodied as LSTMs. This approach is depicted in Figure 3.7. Due to the nature of Poisson distribution, this approach can not produce more than one future step (action, time) at each time step. Notably, both efforts [100, 105] resort to a learning fashion for the hyper-parameter λ.
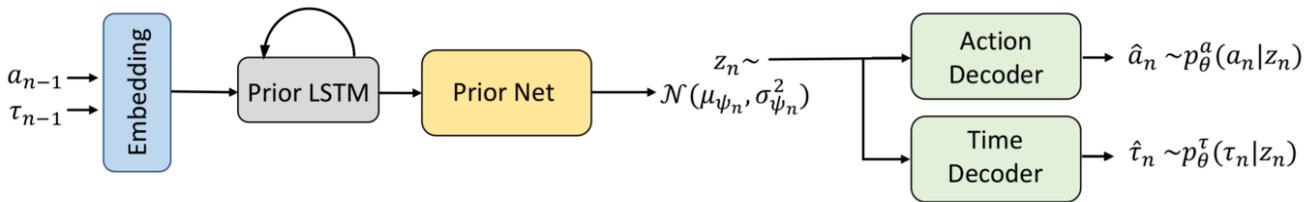
Figure 3.7: Predicting the next-step action and interval time in [105]. Past action and time pairs (a_{n-1}, τ_{n-1}) are used to stochastically produce the future action-time pair from sampled latent variable z_n as well as the Action/Time Decoders. Figure reproduced with permission from [105]

.

Figure 3.8: Model introduced in [2] for action and time prediction via regression. Notice the observed input action is pre-processed into a one-hot encoding and the input time is framed as a real value scalar. Subsequent processing alternates use of fully connected (FC) layers and Gated Recurrent Units (GRUs). The whole scope of the future requires the recursive inference. Figure reproduced with permission from [2].
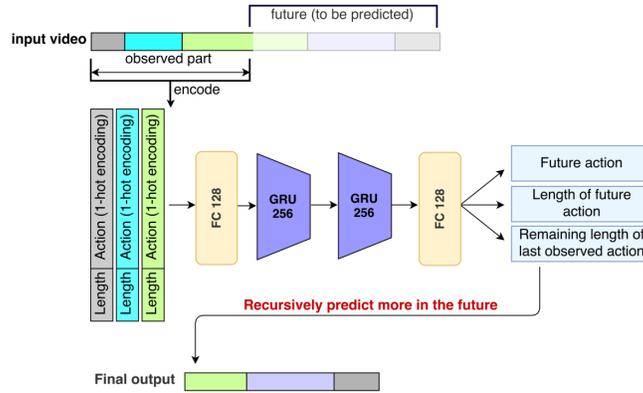
Though the Poisson process has been known for modelling time arrival events, not all researchers adopt it. An alternative research direction modeled time as a real-valued output directly from regression models, e.g., using an MLP to generate the output [1, 2, 30, 61, 101] as shown in Figure 3.8: For any next-step action, the "length of future action" and "remaining length of current action" will be regressed from fully-connected layers. In comparison to previous Poisson methods [100, 105], these approaches choose deterministic time models rather than the Poisson distribution and therefore their outputs for a given input will be fixed. An exception to this general trend instead assumes a Gaussian prior and approximates the mean and variance via learning [1].

A third way to model temporal data is to treat time in terms of discrete categorical values. Along these lines, research on recommendation systems found that modelling time with a categorical representation and optimizing with cross-entropy consistently outperformed the use of real valued modelling [88]. It is worth noticing that the range of the discrete categorical variable normally needs to be defined beforehand. For example, in image classifications, the ImageNet dataset defines 1000 categories and thus its models (i.e., AlexNet) would output probability score vectors $\in \mathcal{R}^{1000}$. For discrete times, the maximum time index is often chosen as the time range value [110].

To clarify the variability of the design choice for time representation, especially in video understanding tasks, a systematic study of the aforementioned choices was conducted [110]. This study grouped time representations into three categories, continuous, discrete and hybrid; see Table 3.1.

In empirical comparison, there were three discrete methods (One-in-many Classifier, Binary Classifier and Heuristic Heatmap), three continuous methods (Direct Regression, Gaussian Distribution and Weibull Distribution) and one hybrid method (Gaussian Mixture of Heuristic Heatmap). These time rep-

Group	Method	Detail
Discrete	One-in-many Classification	Quantize time in Δmax discrete bins and use classified index as results.
	Binary Classifiers	Δmax independent binary classifiers for each time index.
	Heuristic Heat-Map	Generate 1d heat-map and choose the maximum position in heatmap.
Continuous	Direct Regression	Real value regression.
	Gaussian Distribution	Estimate time as parameterized Normal $N(t|\mu, \sigma)$ distribution.
	Weibull Distribution	Estimate tme as parameterized Weibull $p(t|\alpha, \beta) = 1 - e^{-(t/\alpha)^\beta}$.
Hybrid	Gaussian Mixture	Gaussian Mixture of heat-maps. A hybrid model.

Table 3.1: Summary of time representations. Δ_{max} stands for the maximum time range in datasets.

Model	EPA[%]	TTEE[s]	MS
One-in-many Classifier	71.78	2.66	N/A
Binary Classifier	75.90	3.21	N/A
Direct Regression	71.11	3.81	N/A
Gaussian Distribution	66.77	5.77	10.91
Weibull Distribution	55.66	8.31	3.20
Heuristic Heatmap	76.15	2.64	3.88
Gaussian Mixture of Heuristic Heatmap	78.70	2.04	2.69

Table 3.2: Comparison results of various time methods on BDD100K dataset. For evaluation metrics, higher Event Prediction Accuracy (EPA) means better performance, whereas lower Time-to-Event Error (TTEE) and Model Surprise (MS) mean better performance. Table reproduced with permission from [110].

resentations are evaluated on the BDD100K car stop anticipation dataset [175] that measures the Event Prediction Accuracy (EPA), Time-To-Event-Error (TTEE) and Model Surprise (MS). In their results (shown in Table 3.2), almost all continuous time modelling methods fall short of achieving comparable performance with discrete methods. Moreover, the hybrid model performs the best. This observation on video understanding agrees with the experimental results from other related fields [88].

One possible explanation to validate the above phenomenon could be the effect of temporal scope prior, Δ_{max}, enforced by discrete models. Continuous methods generate outputs in the real positive number realm, $\mathcal{R}^+$, while discrete methods pre-confine the output range to $\mathcal{R}^{\lambda \Delta_{max}}$ (λ a scaling factor). This prior implicitly informs the learned prediction model about the reasonable temporal output range, whereas the continuous counterpart tries to cover the whole positive value realm. Another potential reason for the observed pattern of results could be the time data format. Most such work (e.g., [2, 61, 105]) assumes the minimum unit of time is in the unit of seconds, which corresponds to the data formatting common in adopted datasets (e.g., Breakfast dataset [76]). Discrete time format matches the original scheme naturally while continuous ones incur unnecessary decimals. Recent work succeeds the discrete time scheme and trains a conditional GAN model based on Gumbel discrete sampling to jointly enhance the accuracy and diversity in both future action semantics and times [183]; see Figure 3.9. Previous efforts tended to be challenged in producing realistic (accurate) and diverse predictions.

In summary, approaches to joint prediction of actions and their times can be categorized along three dimensions, as follows.

Difference in Inputs: Similar to the discussion on input data use in Section 3.2, research in time prediction also can be categorized according to the input data type. Some work aims at learning the mapping from high-level information, extracted from raw observations, to the future. These efforts usually depend on extracted/estimated activity labels and segmented temporal duration as input [1,2,61,105,183]. Other work relies on raw video inputs [101,110].

Difference in stochastic process: The majority of work relies on a stochastic process for future actions and times. Various approaches employ a Poisson [100,105] or Gaussian distributions [1] or a generative adversarial network [183]. A few outliers deterministically produce their anticipations [2,88].
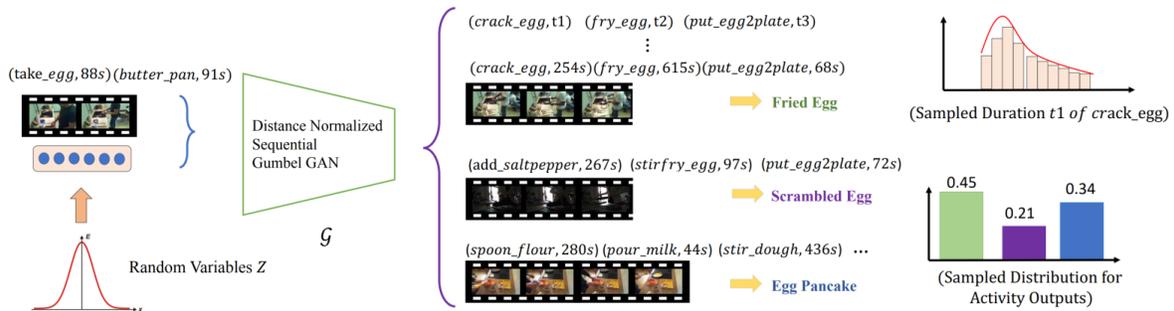
Figure 3.9: Method introduced in [183] that used a sequential modeller (Distance Normalized Gumbel GAN) to produce the whole sequence of subsequent actions and times. Multi-modality is well supported for both action categories and time durations. Figure reproduced with permission from [183].

Difference in prediction horizon: Some work advocates for one-step future predictions. RNNs have been used to recursively generate next step (action, time) pairs (e.g., [2]). To obtain the whole scope of the future, results of next-step prediction are reused as input for the next-next. This procedure repeats until it reaches the end of sequence (e.g., [30,101,105,110]). Alternatively, other approaches generate the rest of the sequence in one shot (e.g., [1,183]). Finally, work has made one-shot time independent predictions, so that predictions on remote time indices do not rely on previous ones [61]. They achieve this goal by making predictions conditioned on the desired prediction index as an extra input. The evident downside of their work is the limited prediction scope.

3.4 Ego-Centric Action Prediction

First person video analysis has attracted increasing attention due to the rich information contained in egocentric visual data and easier access to wearable recording devices. As seen in Figure 3.10, ego-centric filming exhibits much closer viewing angles of objects, environments and interactions of daily activity. However, one critical difference from the third person view dataset (e.g., UCF-101 [145]) is the lack of notable motion of actors, i.e., only the hand motion is typically visible. Moreover, since the camera moves along with its carrier, there usually exists large background motion as well as jittering in video frames.

Figure 3.10: Examples of Ego-centric activity footage.

Missing actors in activity videos makes classic human action analysis methodologies not naturally applicable and larger view-point changes as well as camera jittering makes motion feature extraction (e.g., optical flow) very noisy. These obstacles incentivise researchers to tackle ego-centric data in a slightly different way.

In the following, we group related research work into two main divisions: **1.** work that actively explores non-traditional visual cues (i.e., hand motion, gaze salience, object trajectory, etc.) that are uniquely salient in ego-view data; **2.** work that follows regular techniques and treats ego-centric data in a similar way as third-person view data. In this approach, even though videos are filmed under ego-centric view, performed actions can be inferred without direct observation of large portions of the actor. We first discuss the former division, unique features, from the perspectives of gaze, hand and object.

Gaze Information: One beneficial information source hidden in first person video is gaze information. Gaze moving on/off of a certain object can closely correspond to the occurrence/ending of a particular action. Much work has enjoyed the incorporation of such information (e.g., [85, 138, 148]).

Yet, gaze localization does not come freely. One way to obtain gaze information is the 3D reconstruction of actors and then inferring the gaze direction from head direction and body poses [148]. In that work, the authors virtually stabilized input first person video and reconstructed 3D body pose by applying cylindrical projection using existing tools [85]. This processing provided labels for the location, orientation,

and velocity of actors with pixel level precision, as seen in Figure 3.11 (a). The inferred gaze direction and body pose together indicate the potential social interactions among players and thus contribute to anticipating future behaviors.
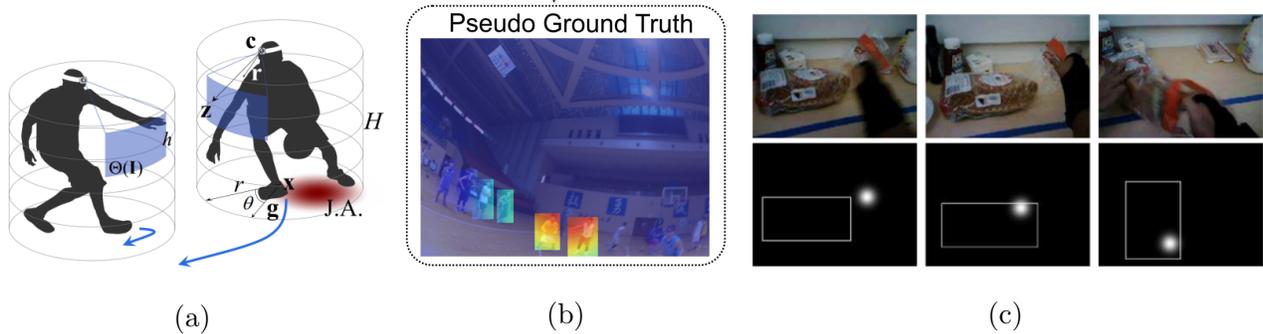
Figure 3.11: (a) Examples of 3D reconstruction of stablized basketball players and the recovered gaze direction **z** from [148]. (b) Example of next player intention groundtruth generated from ego-supervised prior knowledge with high probability score denoted as red, while low score is blue. (c) Gaze-events defined from [138] where object target detection (white rectangles) **overlaps** with gaze fixations (white circle). Subfigures (a), (b) and (c) reproduced with permission from [148], [6] and [138], respectively.

Another way to gain gaze information is through manual annotation and/or use of additional instrumentation. Though expensive, it is easily doable through eye-tracking devices. Assuming decent gaze estimation is given, an effort proposed a model to focus on video frames where target objects and gaze estimations coincide [138] (shown in Figure 3.11 (c)). They paid special attention to gaze-events and distilled attention values from event-fired frames to assist action prediction.

Some other work that considers gaze as an important visual cue focused efforts on predicting future gaze fixations from observed frames [180]. Their framework proceeded in a similar way to video frame prediction but output binary gaze heat-maps. However, the authors did not show any application of the predicted gaze on other related problems.

Hand Information: In first person videos, hand movement typically is the major human visible action. Human-object interactions often consist of a movement part and one or more object parts, with "take bowl" as an example. First person videos capture complex hand movements during a rich set of interactions, thus providing a powerful workhorse for studying the connection between hand actions and future representation.

In imitation of the acquisition of gaze information, hand motion can be either annotated or online estimated. For example, an extant hand shape detector [185] can be applied to obtain a hand mask across frames [138]. The use of the estimated hand mask is straightforward: Overlaying the hand mask on RGB input frames. In other recent work, it is proposed to use the hand movement trajectory as an output

regularizer [94]. The authors enforced the learned network model to jointly predict future actions, future hand trajectories and interactive hotspots (regions where interaction would happen with high likelihood). In an ablation study, it was shown that the three related tasks were complementary to each other.

Object Information: Objects that are interacted with are another potentially important source of information for predicting future actions. While such information can be generally useful for action prediction, it can be of particular importance in first-person video where the movements of the actor may be obscured.

An early work along these lines investigated the relevance of egocentric object trajectories in the task of next-active-object prediction [27]. Provided that an object detector/tracker is available, they proposed to analyze object trajectories observed in a small temporal window to detect next-active-objects before the object-interaction is actually started.

A more recent work incorporated information on the spatial location of objects into ego-centric action prediction [29]. The authors adopted a pretrained deep object detector to localize objects and merged their deep features with another two feature modalities (i.e., RGB and optical flow) captured globally from the entire frame. However, their major inference engine is a novel Rolling-Unrolling (RU) LSTM system and object features are simply used as input. In an ablation study, they found that inclusion of the object information helped achieve much better results compared to neglecting that information.

As mentioned above, there are other efforts that treat ego-data no differently from regular third person data. One such approach processed ego-centric video with a standard deep neural network action recognition model pretrained on third person video datasets [19, 164], but now fine-tuned for ego-action future prediction assisted by top-loss [28]. Other research made use of standard third person analysis tools to anticipate events that might occur in different places than those currently in view [8, 125]. For example, the current view is of the living room, while the future activity will occur in the kitchen. In response, they proposed to merge 3D point cloud location information obtained from a SLAM system with interactive objects and high-level scene labels to do the predicting

Another interesting use of ego-centric data is exploiting the prior knowledge behind the ego-centric data for unsuperivsed learning [6]. The authors researched predicting the next cooperative player with regard to the ego-centric viewer. Instead of manually annotating the groundtruth target, they generated pseudo groundtruth in an unsupervised fashion by using ego-centric assumptions; see Figure 3.11 (b).

3.5 Datasets and Performance

In this section, we describe datasets, evaluation metrics and performance in evaluation of the majority of future action prediction approaches discussed above.

	Year	No. Videos	No. Actions	Avg. Len	Domain	Data-Modality
BDD100K [175]	2020	100,000	3	40s	Car-Stop	RGB
EGTEA [23]	2019	39,596	149	-	Human Interaction	RGB, HandMask, Gaze
PIE [121]	2019	56,000	60	-	Action	RGB, Skeleton
Epic-Kitchen [19]	2018	272	125	3s	Activity, HOI	RGB
MultiTHUMOS [174]	2017	400	65	4.03s	Activity	RGB
JAAD [73]	2017	346	11	6.39s	Action	RGB, BB
CAD120 [72]	2013	120	20	-	Action, HOI	RGB+D, Skeleton
50 Salads [83]	2012	50	17	-	Action-traffic	RGB, Temporal-Seg
MPII-Cooking [127]	2012	44	65	600s	Activity-cooking	RGB
VIRAT [113]	2011	59	10	219	Activity	RGB+D, Skeleton
UCI-OPPORTUNITY [126]	2010	10h	5	-	Action, HOI	RGB, Sensor
Human-Object-Interaction [38]	2009	10	6	-	Action	RGB

Table 3.3: Summary of datasets for future action prediction.

Epic-Kitchen [19] is a large first-person video dataset, which is captured by 32 subjects in 32 different kitchens. The videos in this dataset contain daily activities of the subjects, i.e., no scripts are provided to instruct the subjects. This makes this dataset very natural and challenging. There are 272 training videos, which are captured by 28 subjects. Each video contains multiple action segments, which are categorized into 125 classes. Since the annotations of the testing videos are not available, the training videos are used to perform cross-validation for evaluation. Specifically, the training videos are randomly split into 7 splits, each containing videos of 4 subjects.

50 Salads [83] contains 50 videos which are performed by 25 subjects. Each subject is preparing two mixed salads. There are 17 fine-grained action classes. Researchers often perform 5-fold cross-validation for evaluation using the splits provided by the authors. There are a total of 50 videos that are filmed over 100 hours.

Breakfast [75] contains 1,712 videos of 52 different actors making breakfast. Overall, there are 48 fine-grained action classes and about 6 action instances for each video. The average duration of the videos is 2.3 minutes and the longest video is 10 minutes.

MultiTHUMOS [174] is a challenging dataset for action recognition, containing 400 videos of 65 different actions. On average, there are 10.5 action class labels per video and 1.5 actions per frame and thus it is suitable for long-term activity anticipations.

BDD100K [175] consists of 100,000 driving video sequences each 40 seconds of length, accompanied with basic sensory data such as GPS, velocity or acceleration. In total, there are 31k vehicle stopping and 21k not-stopping sequences for training, and 4.6k stopping and 3.1k non-stopping sequences for evaluation.

VIRAT [113] is a collection of approximately 25 hours of surveillance video taken from various scenes, with an average of 1.6 hours per scene. Multiple HD cameras are used to capture at 1080p or 720p at rates between 25 to 30 Hz. The view angles of cameras towards dominant ground planes ranged between 20 and 50 degrees by stationing cameras mostly at the top of buildings to record a large number of event instances across area, while avoiding occlusion as much as possible.

PIE [121] comprises 1,842 pedestrian tracks captured using an on-board monocular camera, while driving in urban environments with various street structures and crowd densities. Overall, the ratio of pedestrian non-crossing to crossing events is 2.5 to 1. All video sequences are collected during daylight under clear weather conditions. The videos are continuous, allowing us to observe the pedestrians from the moment they appear in the scene until they go out of the field of view of the camera.

MPII-Cooking [127] contains 44 instances of cooking activity, which are continuously recorded in a realistic setting. Predictable high level activities are about preparing 14 kinds of dishes, including:

making a sandwich, making a pizza, and making an omelet, etc. There are overall 65 different actionlets as building blocks shared among various cooking activities, such as cut, pour, shake, and peel.

GTEA Gaze [23] contains 17 sequences of meal preparation activities performed by 14 different subjects, with the spatial resolution of 640 & $\times$ 480. It contains the subjects' gaze location in each frame and the corresponding activity labels. Its extension, **GTEA Gaze++** [22], contains 37 sequences performed by 6 subjects of preparing 7 types of meals. Their latest collection, EGTEA [90], comes with 10, 321 action instances from 19 verb, 53 noun and 106 action classes.

UCI-OPPORTUNITY [126] was created in a sensor-rich environment for the machine recognition of human activities. They deployed 72 sensors of 10 modalities in 15 wireless and wired networked sensor systems in the environment, on the objects, and on the human body. The data are acquired from 12 subjects performing morning activities, yielding over 25 hours of sensor data. It contains five high level predictable activities (Relaxing, Coffee time, Early Morning, Cleanup, Sandwich time), 13 low level actionlets (e.g., lock, stir, open, release), and 23 interactive objects (e.g., bread, table, glass).

CAD120 [72] has 120 RGB-D videos of four different subjects performing 10 high-level activities. The data is annotated with object affordance and sub-activity labels and includes groundtruth object categories, tracked object bounding boxes and human skeletons. The set of high-level activities are making cereal, taking medicine, stacking objects, unstacking objects, microwaving food, picking objects, cleaning objects, taking food, arranging objects, having a meal. The set of sub-activity labels are reaching, moving, pouring, eating, drinking, opening, placing, closing, scrubbing, null and the set of affordance labels are reachable, movable, pourable, pourto, containable, drinkable, openable, placeable, closable, scrubbable, stationary.

HOI [38] is the Maryland Human-Object Interactions dataset, which consists of six annotated activities: answering a phone call, making a phone call, drinking water, lighting a flash, pouring water into container and spraying. These activities have about three to five action units each. Constituent action units share similar human movements: 1) reaching for an object of interest, 2) grasping the object, 3) manipulating the object, and 4) putting back the object. For each activity, there are 8 to 10 video samples.

Evaluations for future action prediction can be grouped into two division: **1.** Reporting the mean average precision of future action label classifications at a single pre-planned future time horizon; **2.** Reporting the mean average precision over a range of time horizons. The former is widely adopted for activity singleton or sequence anticipations, while the latter for joint activity and time duration anticipations.

In this report, we summarize dataset performances that are frequently and extensively investigated in recent years regarding the above two divisions. Additional results can be found in the original papers. For

the first division, we showcase the next activity anticipation results on the Epic-Kitchen dataset [19]. In its standard setting, the prediction horizon is fixed as the future **1s**. Epic-Kitchen dataset has complete annotations for fine-grained sub-action denoted as **Verb**, interactive objects as **Noun** and high-level action category as **Action**. Accuracies for all three attributes will be reported. Two sets of accuracy, namely top-1 that only accept the best prediction for evaluation and top-5 that would accept 5 best guesses as final results, are listed to shed light on the robustness of their models; see Table 3.4.

Table 3.4: Future action prediction accuracy on future 1 second for Epic-Kitchen. "set1" stands for the test set whose scenes appear in the training set, whereas "set2" for the test set that does not share scenes with the training set.

	Method	Top/Top5 Accuracy		
		Verb	Noun	Action
set1	2S-CNN [19]	29.76/76.03	15.15/38.65	4.32/15.21
	TSN [19]	31.81/76.56	16.22/42.15	6.00/18.21
	TSN+MCE [106]	27.92/73.59	16.09/38.32	10.76/25.28
	Trans R(2+1)D [106]	30.74/76.21	16.47/42.72	8.74/25.44
	RULSTM [29]	33.04/79.55	22.78/50.95	14.39/33.73
	FHOI [94]	36.25/79.15	23.83/51.98	15.42/34.29
set2	2S-CNN [19]	25.23/68.66	9.97/27.38	2.29/9.35
	TSN [19]	25.30/68.32	10.41/29.50	2.39/9.63
	TSN+MCE [106]	21.27/63.66	9.90/25.50	5.57/25.28
	Trans R(2+1)D [106]	28.37/69.96	12.43/32.20	7.24/19.29
	RULSTM [29]	27.01/69.55	15.19/34.28	8.16/21.20
	FHOI [94]	29.87/71.77	16.80/38.96	9.94/23.69

For the second division, we follow the Mean-over-Class (MoC) evaluation metrics to measure the performance of joint activity category and duration anticipation on two datasets, Breakfast [75] and 50Salads [83]; see Tables 3.5 and 3.6 for Breakfast and Tables 3.7 and 3.8 for 50Salads. The evaluation method, MoC, represents the average accuracy over the entire prediction horizon.

For singleton activity prediction, as seen in Table 3.4, the current performance is far from satisfactory and thus definitely deserves more effort. For the joint anticipation of activity and time duration, the performance produced from methods that take groundtruth (action, time) observations as input achieves decent results even for remote time indices (e.g., 75% accuracy for 50% future video horizon predictions for the Breakfast dataset [30], as seen in Table 3.7). However, given (action, time) inputs from off-the-self algorithms (Table 3.6 and 3.8), which do not guarantee noisy-free estimation, performance on both datasets drops dramatically. This observation reflects the fact that the current bottle-neck is the accuracy of action detection and segmentation methods more so than the prediction algorithms.

Observation	20%				30%			
Prediction	10%	20%	30%	50%	10%	20%	30%	50%
RNN [2]	0.6035	0.5044	0.4528	0.4042	0.6145	0.5025	0.4490	0.4175
CNN [2]	0.5797	0.4912	0.4403	0.3926	0.6032	0.5014	0.4518	0.4051
TOS-Dense [61]	0.6446	0.5627	0.5015	0.4399	0.6595	0.5594	0.4914	0.4423
R-HMM (**Avg**) [1]	0.5039	0.4171	0.3779	0.3278	0.5125	0.4294	0.3833	0.3307
R-HMM (**Max**) [1]	0.7884	0.7284	0.6629	0.6345	0.8200	0.7283	0.6913	0.6239
NDR-GAN (**Avg**) [183]	0.7222	0.6240	0.5622	0.4595	0.7414	0.7132	0.6530	0.5238
NDR-GAN (**Max**) [183]	0.8208	0.7059	0.6851	0.6406	0.8336	0.7685	0.7213	0.6406
NeualMemory [30]	0.8720	0.8524	0.8102	0.7547	0.8790	0.8279	0.8210	0.7630

Table 3.5: Breakfast Dataset results of dense anticipation mean over classes (MoC) accuracy with groundtruth (action, time) observations as inputs. Avg stands for averaged results across 16 samplings, while Max stands for taking the best result among 16 samples.

Observation	20%				30%			
Prediction	10%	20%	30%	50%	10%	20%	30%	50%
RNN [2]	0.1811	0.1720	0.1594	0.1581	0.2164	0.2002	0.1973	0.1921
CNN [2]	0.1790	0.1635	0.1537	0.1454	0.2244	0.2012	0.1969	0.1876
TOS-Dense [61]	0.1841	0.1721	0.1642	0.1584	0.2275	0.2044	0.1964	0.1975
R-HMM (**mode**) [1]	0.1671	0.1540	0.1447	0.1420	0.2073	0.1827	0.1842	0.1686
TCN-Cycle [21]	0.2588	0.2342	0.2242	0.2154	0.2966	0.2737	0.2558	0.2520
Attn-GRU [111]	0.2303	0.2228	0.2200	0.2085	0.2650	0.2500	0.2408	0.2361

Table 3.6: Breakfast Dataset results of dense anticipation mean over classes (MoC) accuracy **without** groundtruth (action, time) observations as inputs.

Observation	20%				30%			
Prediction	10%	20%	30%	50%	10%	20%	30%	50%
RNN [2]	0.4230	0.3119	0.2522	0.1682	0.4419	0.2951	0.1996	0.1038
CNN [2]	0.3608	0.2762	0.2143	0.1548	0.3736	0.2478	0.2078	0.1405
TOS-Dense [61]	0.4512	0.3323	0.2759	0.1727	**0.4640**	0.3480	0.2524	0.1384
R-HMM (**Avg**) [1]	0.3495	0.2805	0.2408	0.1541	0.3315	0.2465	0.1884	0.1434
R-HMM (**Max**) [1]	0.7489	0.5875	0.4607	0.3571	0.6739	0.5237	0.4673	0.3664
NDR-GAN (**Avg**) [183]	0.4663	0.3562	0.3191	0.2137	0.4613	0.3637	0.3310	0.1945
NDR-GAN (**Max**) [183]	0.5150	0.3845	0.3606	0.2762	0.5079	0.4754	0.3783	0.2908
NeualMemory [30]	0.6996	0.6433	0.6271	0.5216	0.6810	0.6229	0.6118	0.5667

Table 3.7: Results for the 50Salads dataset of dense anticipation mean over classes (MoC) accuracy with groundtruth (action, time) observations as inputs.

Observation	20%				30%			
Prediction	10%	20%	30%	50%	10%	20%	30%	50%
RNN [2]	0.3006	0.2543	0.1874	0.1349	0.2164	0.2002	0.1973	0.1921
CNN [2]	0.2124	0.1903	0.1598	0.0987	0.2914	0.2014	0.1746	0.1086
TOS-Dense [61]	0.3251	0.2761	0.2126	0.1599	0.33512	0.2705	0.2205	0.1559
R-HMM (**mode**) [1]	0.2486	0.2237	0.1988	0.1282	0.2910	0.2050	0.1528	0.1231
TCN-Cycle [21]	0.3476	0.2841	0.2182	0.1525	0.3439	0.2370	0.1895	0.1589
Attn-GRU [111]	0.3932	0.3139	0.2701	0.2388	0.4173	0.3273	0.3144	0.2639
ACC-Grammar [117]	0.3950	0.3320	0.2590	0.2120	0.3950	0.3150	0.2640	0.1980

Table 3.8: Results for the 50Salads dataset of dense anticipation mean over classes (MoC) accuracy **without** groundtruth (action, time) observations as inputs.

Chapter 4

Conclusion and Future Work

4.1 Current trends

Throughout this survey, various video predictive tasks have been discussed, among which early action recognition and future action prediction have been studied in a detailed manner. Even though recent efforts have advanced the performance of both tasks greatly, in this section we summarize current status and discuss the potential urgent challenges awaiting to be solved.

4.1.1 Early action recognition

Current trending approaches on early action recognition greatly rely on success from the video action recognition area, where large scale datasets and powerful video discriminative feature representation are frequently advanced via deep learning. As examples: Recent progress comes from use of enhanced video features (e.g., [70, 182] Temporal segment network [164], [165] ResNet [46] and [139] Inception network [152]). Among this progress, the mutual information based (Section 3.2.2) and the propagation based (Section 3.2.3) are the front runners in terms of performance on recent benchmark datasets. These two approaches share the intuition that partial information needs to be augmented, yet they follow different methodologies to achieve that end.

Mutual information based approaches focus on reconstructing a complete video representation through feature matching or knowledge distillation. In these approaches, partial features are mapped to global features, mostly with feature-wise Euclidean distance measuring the feature similarity, with little or no explicit temporal modeling. In contrast, propagation based approaches focus on the temporal evolution of videos and endeavor to recover useful features at each time step. Some classic temporal sequential

models (e.g., Kalman filtering and Markov Decision process) are frequently revisited in these approaches; see [179, 182].

For early action recognition, a remaining demand is extending study to un-constrained environments. Popular datasets (e.g., UCF-101 [145]) in existing research are known to be carefully trimmed and controlled. Though they are collected from real-life video sources and called "in the wild", the contained motion is almost always clean, visible and smooth. Such quality is not always guaranteed in many real-life environments. Moreover, the fact that early action recognition still performs poorly in the BIT dataset indicates that most current methods are sensitive to the dataset, e.g., between class variation is much greater in UCF-101 compared to BIT.

Even though some research has revealed that certain action categories are relatively easy to recognize early [70], it would be favorable to understand the reason behind the observed patterns. For datasets that can be readily recognized early on (e.g., UT-Interaction), the decisive factor for the success is still unknown. Also as noted by [184], for more complex daily actions that would include compound reasoning structure, i.e., "trying to pour water into a glass, but missing so it spills next to it" from the Something2something dataset, the accuracy is extremely unsatisfactory (around 10% accuracy), indicating that current high accuracy on most examined datasets does not equal to the desired reasoning ability in computational intelligence.

Another critical issue is the current evaluation procedure for early action recognition. As mentioned in Section 3.1, action videos are often being divided into $k \in (10\%, 100\%)$ progress levels and the recognition accuracy with regard to each level will be reported. However, such a split design may be too coarse to allow for understanding of real actions that unfold more incrementally. For some actions that can be completely executed within 10% of frames, a high recognition accuracy is misleading. Furthermore, little has been done to consider the relative importance of early recognition as a function of action (e.g., elder people falling requires earlier recognition to be useful compared to many other actions).

4.1.2 Future Action Prediction

Mainstream approaches for future action prediction either use temporal sequential tools for building the causal relation between past and long-term future at a high-level (e.g., semantics and/or time durations), or make use of visual features (e.g., objects, eye gazes, hand motions, poses, etc.) extracted from recently observed frames for anticipating the immediate next action label. In the first research line, the Recurrent Neural Network (RNN) as well as its hybrids with certain traditional sequence models (e.g., RNN-HMM [1] or Poission-LSTM [105]) are the major workhorse, as seen in recent efforts. The high level semantics are often preprocessed with off-the-shelf algorithms and these approaches, in most cases, resorted to a data-

driven methodology to learn parameters for both RNNs and classic models.

In the second direction (i.e., more direct use of visual features), the focus is employing multi-modal data with the most effectiveness. Typically, the fusion of distinct information is achieved through one of the following strategies: stacking them together as a single input (e.g., [29]); designing a multi-task learning framework with extra data modality as regularizers (e.g., [94]); fusing them using a hierarchical structure (e.g., [56, 123]). As means to produce the future action label, researchers often resorted to the recently advanced deep ConvNet features (e.g., I3D features [11]) and learned a classifier from the training data.

For joint activity and time duration prediction, it is still an open question as to how to represent times most effectively. Modelling time as a discrete variable seems to perform well, yet it fails when any future time duration exceeds its pre-defined scope. In contrast, the continuous representation would either be too coarse (e.g., real value regression) or constrained into a uni-modal distribution (e.g., Gaussian), which naturally lacks the ability to model multi-modality of times (e.g., *boiling_egg* can take longer time for a hard boiled egg as well as shorter time for a soft boiled one).

Currently, the prediction of activities often happens in a narrow spatial scope, such as kitchen areas or cooking stations. Yet, daily activities often take place with a transition of locations, i.e., fetching objects in another room can incur a scene change. Even for the ambitious Epic-Kitchen dataset [19], the default assumption is that all objects relevant for prediction are visible in the current frame. Indeed, there seems to be only one effort that considers prediction across multiple locations [125] and the dataset has yet to be shared.

As discussed in Section 3.2, in long-term future action predictions, researchers typically resort to the mapping between the high-level semantics [86] (e.g., model a mapping from a sequence of observed activities to a sequence of future activity labels). Indeed, in some cases researchers have gone to the extreme of taking groundtruth labels for the observed portion of the video as input. This setting is reasonable in that an extended video clip contains many frames, which prohibit online low-level feature computation as well as storage, and high-level semantics summarize long videos in a manageable and meaningful way. Nonetheless, this scheme ignores the rich visual elements that have been demonstrated to be worth exploiting in other areas and which are worth further consideration in long-term predictions. Moreover, approaches that employ groundtruth labels as input avoid the critical issue of how to extract the labels video in the first place.

4.2 Future Work

In this final section, several directions for future work are suggested.

Improved interpretability of algorithms and representations is necessary to truly understand the early action recognition and future action prediction tasks and their enabling mechanisms. From the perspective of understanding the learning methods, there is a demand to further research *what has been learned* by proposed solutions. As an example: In student-teacher learning [165], it would be beneficial to monitor and understand the teaching process. In the example of deep Kalman feature propagation [182], answers are needed for questions such as how does learning based Kalman filtering affect the error accumulation and what are the motion kernels learned for feature propagation (e.g., are they mimicking a warping function)? Both analytic (e.g., [41]) and visualization methodologies (e.g., [24]) have potential for shedding light on these questions.

Integrating feature extraction with prediction is an urgent concern for most high-level based approaches. It has been seen that a decent semantic estimation of input partial observations plays a critical role in making satisfactory long-term prediction (e.g., see Table 3.5 and Table 3.7), while a noisy estimation deteriorates the performance remarkably (e.g., see Table 3.6 and Table 3.8). To tackle this discrepancy, it might be worth integrating the future prediction into the semantic extraction, with an objective to improve the robustness of future prediction with noisy inputs.

Better modelling of long-term correlation is necessary to achieve scalable prediction of activities. The cue that can precisely indicate the future could possibly hide in details that are far away from current observations, and thus it is essential to keep logging historical information and effectively spot the correlations between minor details and future events. Though some efforts have been made for action recognition (e.g., the authors in [166] designed a long-term filter bank to keep a track of remote historical features), its effects on other areas stays unknown and it is likely beneficial to revisit such an ability for future action predictions.

Stronger temporal sequence modelling tools are always favorable to video understanding. The majority of mathematical tools mentioned in Chapter 2 are theoretically grounded and widely adopted. Recently, however, it has been proposed to explore other advances, such as temporal convolution [83], which challenges the use of memory hidden states, or transformers [34], which get rid of the recurrent structure and yet demonstrates superior results, as alternatives to classic sequence modelling tools. These innovate findings break the traditional design where the temporal modelling has to be sequential and mostly Markovian. Such directions are worth further investigation.

Learning efficiency is another concern that closely associates with most recent learning approaches,

since these methods prerequisite sufficient size of datasets and expensive annotations, which incurs huge amounts of human labeling efforts, data cleaning and unavoidable data biases. So, how to learn meaningful information without strong supervision or with noisy labels awaits to be answered. Besides, there is also a need for an adaptive learning framework that can assimilate unseen data and new action categories, since people's daily activities would evolve along with their own development. An active and continuous learning framework would be of great benefit. Unsupervised approaches to representational learning, e.g., as introduced under the topic of predictive video coding in Chapter 1, provide a potentially useful avenue to explore to address these concerns.

Creation of better datasets and experimental protocols is necessary to better validate prediction algorithms. From the perspective of datasets and evaluation metrics, it is worth understanding the specific observation ratios that are practically valuable in real-life conditions with regard to actions, and therefore re-evaluating existing approaches using this new criterion. It is also worth investigating the correlation between the video observation ratio and action evolution ratio of popular datasets, in order to examine the validation of current early recognition evaluation metrics.

Acknowledgements

References

[1] Yazan Abu Farha and Juergen Gall. Uncertainty-aware anticipation of activities. In ICCVW, pages 1197–1204, 2019.

[2] Yazan Abu Farha, Alexander Richard, and Juergen Gall. When will you do what? -Anticipating temporal occurrences of activities. In CVPR, pages 5343–5352, 2018.

[3] Alexandre Alahi, Kratarth Goel, Vignesh Ramanathan, Alexandre Robicquet, Li Fei-Fei, and Silvio Savarese. Social LSTM: Human trajectory prediction in crowded spaces. In CVPR, pages 961–971, 2016.

[4] Sadegh Aliakbarian, Fatemeh Sadat Saleh, Mathieu Salzmann, Lars Petersson, and Stephen Gould. A stochastic conditioning scheme for diverse human motion prediction. In CVPR, pages 5223–5232, 2020.

[5] Mohammad Babaeizadeh, Chelsea Finn, Dumitru Erhan, Roy H Campbell, and Sergey Levine. Stochastic variational video prediction. arXiv preprint arXiv:1710.11252, 2017.

[6] Gedas Bertasius and Jianbo Shi. Using cross-model egosupervision to learn cooperative basketball intention. In ICCVW, pages 2355–2363, 2017.

[7] Hakan Bilen, Basura Fernando, Efstratios Gavves, and Andrea Vedaldi. Action recognition with dynamic image networks. TPAMI, 40(12):2799–2813, 2017.

[8] Syed Zahir Bokhari and Kris M Kitani. Long-term activity forecasting using first-person vision. In ACCV, pages 346–360, 2016.

[9] Michael Calonder, Vincent Lepetit, Christoph Strecha, and Pascal Fua. BRIEF: Binary robust independent elementary features. In ECCV, pages 778–792, 2010.

[10] Yu Cao, Daniel Barrett, Andrei Barbu, Siddharth Narayanaswamy, Haonan Yu, Aaron Michaux, Yuewei Lin, Sven Dickinson, Jeffrey Mark Siskind, and Song Wang. Recognize human activities from partially observed videos. In CVPR, pages 2658–2665, 2013.

[11] Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? a new model and the Kinetics dataset. In CVPR, pages 6299–6308, 2017.

[12] Anirban Chakraborty and Amit K Roy-Chowdhury. Context-aware activity forecasting. In ACCV, pages 21–36, 2014.

[13] Junwen Chen, Wentao Bao, and Yu Kong. Group activity prediction with sequential relational anticipation model. arXiv preprint arXiv:2008.02441, 2020.

REFERENCES

[14] Lei Chen, Jiwen Lu, Zhanjie Song, and Jie Zhou. Part-activated deep reinforcement learning for action prediction. In ECCV, pages 421–436, 2018.

[15] Minmin Chen, Zhixiang Xu, Kilian Weinberger, and Fei Sha. Marginalized denoising autoencoders for domain adaptation. arXiv preprint arXiv:1206.4683, 2012.

[16] Corinna Cortes and Vladimir Vapnik. Support-vector networks. Machine learning, 20(3):273–297, 1995.

[17] Huseyin Coskun, David Joseph Tan, Sailesh Conjeti, Nassir Navab, and Federico Tombari. Human motion analysis with deep metric learning. In ECCV, pages 667–683, 2018.

[18] Navneet Dalal and Bill Triggs. Histograms of oriented gradients for human detection. In CVPR, pages 886–893, 2005.

[19] Dima Damen, Hazel Doughty, Giovanni Maria Farinella, Sanja Fidler, Antonino Furnari, Evangelos Kazakos, Davide Moltisanti, Jonathan Munro, Toby Perrett, Will Price, et al. Scaling egocentric vision: The Epic-Kitchens dataset. In ECCV, pages 720–736, 2018.

[20] Dave Epstein, Boyuan Chen, and Carl Vondrick. Oops! predicting unintentional action in video. In Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pages 919–929, 2020.

[21] Yazan Abu Farha, Qiuhong Ke, Bernt Schiele, and Juergen Gall. Long-term anticipation of activities with cycle consistency. arXiv preprint arXiv:2009.01142, 2020.

[22] Alireza Fathi, Yin Li, and James M Rehg. Learning to recognize daily actions using gaze. In ECCV, pages 314–327, 2012.

[23] Alireza Fathi, Xiaofeng Ren, and James M Rehg. Learning to recognize objects in egocentric activities. In CVPR, pages 3281–3288, 2011.

[24] Christoph Feichtenhofer, Axel Pinz, Richard P Wildes, and Andrew Zisserman. Deep insights into convolutional networks for video recognition. IJCV, 128(2):420–437, 2020.

[25] Basura Fernando and Samitha Herath. Anticipating human actions by correlating past with the future with jaccard similarity measures. arXiv preprint arXiv:2105.12414, 2021.

[26] Katerina Fragkiadaki, Sergey Levine, Panna Felsen, and Jitendra Malik. Recurrent network models for human dynamics. In ICCV, pages 4346–4354, 2015.

[27] Antonino Furnari, Sebastiano Battiato, Kristen Grauman, and Giovanni Maria Farinella. Next-active-object prediction from egocentric videos. Journal of Visual Communication and Image Representation, 49:401–411, 2017.

[28] Antonino Furnari, Sebastiano Battiato, and Giovanni Maria Farinella. Leveraging uncertainty to rethink loss functions and evaluation measures for egocentric action anticipation. In ECCV, pages 389–405, 2018.

[29] Antonino Furnari and Giovanni Maria Farinella. What would you expect? anticipating egocentric actions with rolling-unrolling LSTM and modality attention. In ICCV, pages 6252–6261, 2019.

[30] Harshala Gammulle, Simon Denman, Sridha Sridharan, and Clinton Fookes. Forecasting future action sequences with neural memory networks. arXiv preprint arXiv:1909.09278, 2019.

REFERENCES

[31] Harshala Gammulle, Simon Denman, Sridha Sridharan, and Clinton Fookes. Predicting the future: A jointly learnt model for action anticipation. In ICCV, pages 5562–5571, 2019.

[32] Jiyang Gao, Zhenheng Yang, and Ram Nevatia. RED: Reinforced encoder-decoder networks for action anticipation. arXiv preprint arXiv:1707.04818, 2017.

[33] Stuart Geman and Christine Graffigne. Markov random field image models and their applications to computer vision. In International Congress of Mathematicians, pages 1496–1517, 1986.

[34] Rohit Girdhar, Joao Carreira, Carl Doersch, and Andrew Zisserman. Video action transformer network. In CVPR, pages 244–253, 2019.

[35] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In NIPS, pages 2672–2680, 2014.

[36] Raghav Goyal, Samira Ebrahimi Kahou, Vincent Michalski, Joanna Materzynska, Susanne Westphal, Heuna Kim, Valentin Haenel, Ingo Fruend, Peter Yianilos, Moritz Mueller-Freitag, et al. The "something something" video database for learning and evaluating visual common sense. In ICCV, pages 5843–5851, 2017.

[37] Liang-Yan Gui, Yu-Xiong Wang, Xiaodan Liang, and José MF Moura. Adversarial geometry-aware human motion prediction. In ECCV, pages 786–803, 2018.

[38] Abhinav Gupta, Aniruddha Kembhavi, and Larry S Davis. Observing human-object interactions: Using spatial and functional compatibility for recognition. TPAMI, 31(10):1775–1789, 2009.

[39] Abhinav Gupta, Praveen Srinivasan, Jianbo Shi, and Larry S. Davis. *Understanding videos, constructing plots* Learning a visually grounded storyline model from annotated videos. In CVPR, pages 2012–2019, 2009.

[40] Agrim Gupta, Justin Johnson, Li Fei-Fei, Silvio Savarese, and Alexandre Alahi. Social GAN: Socially acceptable trajectories with generative adversarial networks. In CVPR, pages 2255–2264, 2018.

[41] Isma Hadji and Richard P Wildes. A spatiotemporal oriented energy network for dynamic texture recognition. In ICCV, pages 3066–3074, 2017.

[42] Raffay Hamid, Siddhartha Maddi, Amos Johnson, Aaron Bobick, Irfan Essa, and Charles Isbell. A novel sequence representation for unsupervised analysis of human activities. Artificial Intelligence, 173(14):1221–1244, 2009.

[43] William L Hamilton, Rex Ying, and Jure Leskovec. Representation learning on graphs: Methods and applications. arXiv preprint arXiv:1709.05584, 2017.

[44] Tengda Han, Weidi Xie, and Andrew Zisserman. Video representation learning by dense predictive coding. In ICCVW, pages 1483–1492, 2019.

[45] Christopher G Harris, Mike Stephens, et al. A combined corner and edge detector. In Alvey Vision Conference, pages 147–151, 1988.

[46] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In CVPR, pages 770–778, 2016.

[47] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. arXiv preprint arXiv:1503.02531, 2015.

[48] Minh Hoai and Fernando De la Torre. Max-margin early event detectors. IJCV, 107(2):191–202, 2014.

[49] Jingyi Hou, Xinxiao Wu, Ruiqi Wang, Jiebo Luo, and Yunde Jia. Confidence-guided self refinement for action prediction in untrimmed videos. TIP, 29(4):6017–6031, 2020.

[50] Jun-Ting Hsieh, Bingbin Liu, De-An Huang, Li F Fei-Fei, and Juan Carlos Niebles. Learning to decompose and disentangle representations for video prediction. In NIPS, pages 517–526, 2018.

[51] Jian-Fang Hu, Wei-Shi Zheng, Jianhuang Lai, and Jianguo Zhang. Jointly learning heterogeneous features for RGB-D activity recognition. In CVPR, pages 5344–5352, 2015.

[52] Jian-Fang Hu, Wei-Shi Zheng, Lianyang Ma, Gang Wang, Jianhuang Lai, and Jianguo Zhang. Early action prediction by soft regression. TPAMI, 41(11):2568–2583, 2018.

[53] Boris Ivanovic and Marco Pavone. The trajectron: Probabilistic multi-agent trajectory modeling with dynamic spatiotemporal graphs. In ICCV, pages 2375–2384, 2019.

[54] Andrew Jaegle, Oleh Rybkin, Konstantinos G Derpanis, and Kostas Daniilidis. Predicting the future with transformational states. arXiv preprint arXiv:1803.09760, 2018.

[55] Ashesh Jain, Hema S Koppula, Bharad Raghavan, Shane Soh, and Ashutosh Saxena. Car that knows before you do: Anticipating maneuvers via learning temporal driving models. In ICCV, pages 3182–3190, 2015.

[56] Ashesh Jain, Avi Singh, Hema S Koppula, Shane Soh, and Ashutosh Saxena. Recurrent neural networks for driver activity anticipation via sensory-fusion architecture. In ICRA, pages 3118–3125, 2016.

[57] Hueihan Jhuang, Juergen Gall, Silvia Zuffi, Cordelia Schmid, and Michael J Black. Towards understanding action recognition. In ICCV, pages 3192–3199, 2013.

[58] Xiaojie Jin, Huaxin Xiao, Xiaohui Shen, Jimei Yang, Zhe Lin, Yunpeng Chen, Zequn Jie, Jiashi Feng, and Shuicheng Yan. Predicting scene parsing and motion dynamics in the future. In NIPS, pages 6915–6924, 2017.

[59] Soo Min Kang and Richard P Wildes. Review of action recognition and detection methods. arXiv preprint arXiv:1610.06906, 2016.

[60] Andrej Karpathy, George Toderici, Sanketh Shetty, Thomas Leung, Rahul Sukthankar, and Li Fei-Fei. Large-scale video classification with convolutional neural networks. In CVPR, pages 1725–1732, 2014.

[61] Qiuhong Ke, Mario Fritz, and Bernt Schiele. Time-conditioned action anticipation in one shot. In CVPR, pages 9925–9934, 2019.

[62] J.F.C. Kingman. Poisson Processes. Oxford Science Publications. Clarendon Press, Oxford, 1993.

[63] Kris M Kitani, Brian D Ziebart, James Andrew Bagnell, and Martial Hebert. Activity forecasting. In ECCV, pages 201–214, 2012.

[64] Yu Kong and Yun Fu. Max-margin action prediction machine. TPAMI, 38(9):1844–1858, 2015.

[65] Yu Kong and Yun Fu. Human action recognition and prediction: A survey. arXiv preprint arXiv:1806.11230, 2018.

[66] Yu Kong and Shangqian Gao. Action prediction from videos via memorizing hard-to-predict samples. In AAAI, pages 7000–7007, 2018.

[67] Yu Kong, Yunde Jia, and Yun Fu. Learning human interaction by interactive phrases. In ECCV, pages 300–313, 2012.

[68] Yu Kong, Dmitry Kit, and Yun Fu. A discriminative model with multiple temporal scales for action prediction. In ECCV, pages 596–611, 2014.

[69] Yu Kong, Zhiqiang Tao, and Yun Fu. Deep sequential context networks for action prediction. In CVPR, pages 1473–1481, 2017.

[70] Yu Kong, Zhiqiang Tao, and Yun Fu. Adversarial action prediction networks. TPAMI, 42(3):539–553, 2018.

[71] Hema S Koppula and Ashutosh Saxena. Anticipating human activities using object affordances for reactive robotic response. TPAMI, 38(1):14–29, 2015.

[72] Hema Swetha Koppula, Rudhir Gupta, and Ashutosh Saxena. Learning human activities and object affordances from RGB-D videos. IJRR, 38(8):951–970, 2013.

[73] Iuliia Kotseruba, Amir Rasouli, and John K Tsotsos. Joint attention in autonomous driving (jaad). arXiv preprint arXiv:1609.04741, 2016.

[74] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. ImageNet classification with deep convolutional neural networks. In NIPS, pages 1097–1105, 2012.

[75] Hilde Kuehne, Ali Arslan, and Thomas Serre. The language of actions: Recovering the syntax and semantics of goal-directed human activities. In CVPR, pages 780–787, 2014.

[76] Hilde Kuehne, Juergen Gall, and Thomas Serre. An end-to-end generative framework for video segmentation and recognition. In WACV, pages 1–8, 2016.

[77] Hildegard Kuehne, Hueihan Jhuang, Estíbaliz Garrote, Tomaso Poggio, and Thomas Serre. HMDB: a large video database for human motion recognition. In ICCV, pages 2556–2563, 2011.

[78] John Lafferty, Andrew McCallum, and Fernando CN Pereira. Conditional random fields: Probabilistic models for segmenting and labeling sequence data. In ICML, pages 282–289, 2001.

[79] Shaofan Lai, Wei-Shi Zheng, Jian-Fang Hu, and Jianguo Zhang. Global-local temporal saliency action prediction. TIP, 27(5):2272–2285, 2017.

[80] Tian Lan, Tsung-Chuan Chen, and Silvio Savarese. A hierarchical representation for future action prediction. In ECCV, pages 689–704, 2014.

[81] Ivan Laptev. On space-time interest points. IJCV, 64(2-3):107–123, 2005.

[82] Kennard Laviers, Gita Sukthankar, Matthew Molineaux, and David W Aha. Improving offensive performance through opponent modeling. Technical report, KNEXUS Research Corporation, 2009.

[83] Colin Lea, Michael D Flynn, Rene Vidal, Austin Reiter, and Gregory D Hager. Temporal convolutional networks for action segmentation and detection. In CVPR, pages 156–165, 2017.

REFERENCES

[84] Alex X Lee, Richard Zhang, Frederik Ebert, Pieter Abbeel, Chelsea Finn, and Sergey Levine. Stochastic adversarial video prediction. arXiv preprint arXiv:1804.01523, 2018.

[85] Vincent Lepetit, Francesc Moreno-Noguer, and Pascal Fua. EPNP: An accurate O(n) solution to the PNP problem. IJCV, 81(2):155, 2009.

[86] Kang Li and Yun Fu. Prediction of human activity by discovering temporal sequence patterns. TPAMI, 36(8):1644–1657, 2014.

[87] Kang Li, Jie Hu, and Yun Fu. Modeling complex temporal composition of actionlets for activity prediction. In ECCV, pages 286–299, 2012.

[88] Yang Li, Nan Du, and Samy Bengio. Time-dependent representation for neural event sequence prediction. arXiv preprint arXiv:1708.00065, 2017.

[89] Yanghao Li, Cuiling Lan, Junliang Xing, Wenjun Zeng, Chunfeng Yuan, and Jiaying Liu. Online human action detection using joint classification-regression recurrent neural networks. In ECCV, pages 203–220, 2016.

[90] Yin Li, Miao Liu, and James M Rehg. In the eye of beholder: Joint learning of gaze and actions in first person video. In ECCV, pages 619–635, 2018.

[91] Junwei Liang, Lu Jiang, Juan Carlos Niebles, Alexander G Hauptmann, and Li Fei-Fei. Peeking into the future: Predicting future person activities and locations in videos. In CVPRW, pages 5725–5734, 2019.

[92] Chunhui Liu, Yueyu Hu, Yanghao Li, Sijie Song, and Jiaying Liu. PKU-MMD: A large scale benchmark for continuous multi-modal human action understanding. arXiv preprint arXiv:1703.07475, 2017.

[93] Jun Liu, Amir Shahroudy, Gang Wang, Ling-Yu Duan, and Alex C Kot. SSNet: scale selection network for online 3D action prediction. In CVPR, pages 8349–8358, 2018.

[94] Miao Liu, Siyu Tang, Yin Li, and James Rehg. Forecasting human object interaction: Joint prediction of motor attention and egocentric activity. arXiv preprint arXiv:1911.10967, 2019.

[95] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott Reed, Cheng-Yang Fu, and Alexander C Berg. SSD: Single shot multibox detector. In ECCV, pages 21–37, 2016.

[96] Ziwei Liu, Raymond A Yeh, Xiaoou Tang, Yiming Liu, and Aseem Agarwala. Video frame synthesis using deep voxel flow. In ICCV, pages 4463–4471, 2017.

[97] Pauline Luc, Camille Couprie, Yann Lecun, and Jakob Verbeek. Predicting future instance segmentation by forecasting convolutional features. In ECCV, pages 584–599, 2018.

[98] Pauline Luc, Natalia Neverova, Camille Couprie, Jakob Verbeek, and Yann LeCun. Predicting deeper into the future of semantic segmentation. In ICCV, pages 648–657, 2017.

[99] Shugao Ma, Leonid Sigal, and Stan Sclaroff. Learning activity progression in LSTMs for activity detection and early detection. In CVPR, pages 1942–1950, 2016.

[100] Tahmida Mahmud, Mahmudul Hasan, Anirban Chakraborty, and Amit K Roy-Chowdhury. A Poisson process model for activity forecasting. In ICIP, pages 3339–3343, 2016.

REFERENCES

[101] Tahmida Mahmud, Mahmudul Hasan, and Amit K Roy-Chowdhury. Joint prediction of activity labels and starting times in untrimmed videos. In ICCV, pages 5773–5782, 2017.

[102] Julieta Martinez, Michael J Black, and Javier Romero. On human motion prediction using recurrent neural networks. In CVPR, pages 2891–2900, 2017.

[103] Joanna Materzynska, Guillaume Berger, Ingo Bax, and Roland Memisevic. The Jester dataset: A large-scale video dataset of human gestures. In ICCVW, pages 2874–2882, 2019.

[104] Michael Mathieu, Camille Couprie, and Yann LeCun. Deep multi-scale video prediction beyond mean square error. arXiv preprint arXiv:1511.05440, 2015.

[105] Nazanin Mehrasa, Akash Abdu Jyothi, Thibaut Durand, Jiawei He, Leonid Sigal, and Greg Mori. A variational auto-encoder model for stochastic point processes. In CVPR, pages 3165–3174, 2019.

[106] Antoine Miech, Ivan Laptev, Josef Sivic, Heng Wang, Lorenzo Torresani, and Du Tran. Leveraging the present to anticipate the future in videos. In CVPRW, pages 2915–2922, 2019.

[107] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. Human-level control through deep reinforcement learning. Nature, 518(7540):529–533, 2015.

[108] Abduallah Mohamed, Kun Qian, Mohamed Elhoseiny, and Christian Claudel. Social-STGCNN: A social spatio-temporal graph convolutional neural network for human trajectory prediction. In CVPR, pages 14424–14432, 2020.

[109] Kevin Murphy, Yair Weiss, and Michael I Jordan. Loopy belief propagation for approximate inference: An empirical study. arXiv preprint arXiv:1301.6725, 2013.

[110] Lukás Neumann, Andrew Zisserman, and Andrea Vedaldi. Future event prediction: If and when. In CVPR, pages 2935–2943, 2019.

[111] Yan Bin Ng and Basura Fernando. Forecasting future action sequences with attention: a new approach to weakly supervised action forecasting. TIP, 29(8):8880–8891, 2020.

[112] Junhyuk Oh, Xiaoxiao Guo, Honglak Lee, Richard L Lewis, and Satinder Singh. Action-conditional video prediction using deep networks in Atari games. In NIPS, pages 2863–2871, 2015.

[113] Sangmin Oh, Anthony Hoogs, Amitha Perera, Naresh Cuntoor, Chia-Chih Chen, Jong Taek Lee, Saurajit Mukherjee, JK Aggarwal, Hyungtae Lee, Larry Davis, et al. A large-scale benchmark dataset for event recognition in surveillance video. In CVPR, pages 3153–3160, 2011.

[114] Marc Oliu, Javier Selva, and Sergio Escalera. Folded recurrent neural networks for future video prediction. In ECCV, pages 716–731, 2018.

[115] Alonso Patron-Perez, Marcin Marszalek, Andrew Zisserman, and Ian Reid. High five: Recognising human interactions in TV shows. In BMVC, pages 1–11, 2010.

[116] Mingtao Pei, Yunde Jia, and Song-Chun Zhu. Parsing video events with goal inference and intent prediction. In ICCV, pages 487–494, 2011.

[117] AJ Piergiovanni, Anelia Angelova, Alexander Toshev, and Michael S Ryoo. Adversarial generative grammars for human activity prediction. In European Conference on Computer Vision, pages 507–523. Springer, 2020.

[118] Jie Qin, Li Liu, Ling Shao, Bingbing Ni, Chen Chen, Fumin Shen, and Yunhong Wang. Binary coding for partial action analysis with limited observation ratios. In CVPR, pages 146–155, 2017.

[119] MarcAurelio Ranzato, Arthur Szlam, Joan Bruna, Michael Mathieu, Ronan Collobert, and Sumit Chopra. Video (language) modeling: a baseline for generative models of natural videos. arXiv preprint arXiv:1412.6604, 2014.

[120] Michalis Raptis and Leonid Sigal. Poselet key-framing: A model for human activity recognition. In CVPR, pages 2650–2657, 2013.

[121] Amir Rasouli, Iuliia Kotseruba, Toni Kunic, and John K Tsotsos. PIE: A large-scale dataset and models for pedestrian intention estimation and trajectory prediction. In ICCV, pages 6262–6271, 2019.

[122] Amir Rasouli, Iuliia Kotseruba, and John K Tsotsos. Are they going to cross? a benchmark dataset and baseline for pedestrian crosswalk behavior. In ICCVW, pages 206–213, 2017.

[123] Amir Rasouli, Iuliia Kotseruba, and John K Tsotsos. Pedestrian action anticipation using contextual feature fusion in stacked RNNs. arXiv preprint arXiv:2005.06582, 2020.

[124] Fahimeh Rezazadegan, Sareh Shirazi, Mahsa Baktashmotlagh, and Larry S Davis. On encoding temporal evolution for real-time action prediction. arXiv preprint arXiv:1709.07894, 2017.

[125] Nicholas Rhinehart and Kris M Kitani. First-person activity forecasting with online inverse reinforcement learning. In ICCV, pages 3696–3705, 2017.

[126] Daniel Roggen, Alberto Calatroni, Mirco Rossi, Thomas Holleczek, Kilian Förster, Gerhard Tröster, Paul Lukowicz, David Bannach, Gerald Pirkl, Alois Ferscha, et al. Collecting complex activity datasets in highly rich networked sensor environments. In International Conference on Network Sensing Systems, pages 233–240, 2010.

[127] Marcus Rohrbach, Sikandar Amin, Mykhaylo Andriluka, and Bernt Schiele. A database for fine grained activity detection of cooking activities. In CVPR, pages 1194–1201, 2012.

[128] Oleh Rybkin, Karl Pertsch, Konstantinos G Derpanis, Kostas Daniilidis, and Andrew Jaegle. Learning what you can do before doing anything. arXiv preprint arXiv:1806.09655, 2018.

[129] Michael S Ryoo. Human activity prediction: Early recognition of ongoing activities from streaming videos. In ICCV, pages 1036–1043, 2011.

[130] Michael S Ryoo, Chia-Chih Chen, JK Aggarwal, and Amit Roy-Chowdhury. An overview of contest on semantic description of human activities. In ICPR, pages 270–285, 2010.

[131] MS Ryoo and JK Aggarwal. UT-Interaction dataset, ICPR contest on semantic description of human activities. cvrc.ece.utexas.edu/SDHA2010/Human_Interaction.html, 2015.

[132] MS Ryoo, Thomas J Fuchs, Lu Xia, Jake K Aggarwal, and Larry Matthies. Robot-centric activity prediction from first-person videos: What will they do to me? In International Conference on Human-Robot Interaction, pages 295–302, 2015.

[133] Mohammad Sadegh Aliakbarian, Fatemeh Sadat Saleh, Mathieu Salzmann, Basura Fernando, Lars Petersson, and Lars Andersson. Encouraging LSTM to anticipate actions very early. In ICCV, pages 280–289, 2017.

REFERENCES

[134] Amir Sadeghian, Vineet Kosaraju, Ali Sadeghian, Noriaki Hirose, Hamid Rezatofighi, and Silvio Savarese. Sophie: An attentive GAN for predicting paths compliant to social and physical constraints. In CVPR, pages 1349–1358, 2019.

[135] Tim Salzmann, Boris Ivanovic, Punarjay Chakravarty, and Marco Pavone. Trajectron++: Multi-agent generative trajectory forecasting with heterogeneous data for control. arXiv preprint arXiv:2001.03093, 2020.

[136] Adam Santoro, David Raposo, David G Barrett, Mateusz Malinowski, Razvan Pascanu, Peter Battaglia, and Timothy Lillicrap. A simple neural network module for relational reasoning. In NIPS, pages 4967–4976, 2017.

[137] Amir Shahroudy, Jun Liu, Tian-Tsong Ng, and Gang Wang. NTU RGB+D: A large scale dataset for 3D human activity analysis. In CVPR, pages 1010–1019, 2016.

[138] Yang Shen, Bingbing Ni, Zefan Li, and Ning Zhuang. Egocentric activity prediction via event modulated attention. In ECCV, pages 197–212, 2018.

[139] Yuge Shi, Basura Fernando, and Richard Hartley. Action anticipation with rbf kernelized feature mapping RNN. In ECCV, pages 301–317, 2018.

[140] Gurkirt Singh, Suman Saha, and Fabio Cuzzolin. Predicting action tubes. In ECCV, pages 106–123, 2018.

[141] Gurkirt Singh, Suman Saha, Michael Sapienza, Philip HS Torr, and Fabio Cuzzolin. Online real-time multiple spatiotemporal action localisation and prediction. In ICCV, pages 3637–3646, 2017.

[142] Josef Sivic and Andrew Zisserman. Efficient visual search of videos cast as text retrieval. TPAMI, 31(4):591–606, 2008.

[143] Khurram Soomro, Haroon Idrees, and Mubarak Shah. Predicting the where and what of actors and actions through online action localization. In CVPR, pages 2648–2657, 2016.

[144] Khurram Soomro, Haroon Idrees, and Mubarak Shah. Online localization and prediction of actions and interactions. TPAMI, 41(2):459–472, 2018.

[145] Khurram Soomro, Amir Roshan Zamir, and Mubarak Shah. UCF101: A dataset of 101 human actions classes from videos in the wild. arXiv preprint arXiv:1212.0402, 2012.

[146] Nitish Srivastava, Elman Mansimov, and Ruslan Salakhudinov. Unsupervised learning of video representations using LSTMs. In ICML, pages 843–852, 2015.

[147] Kun Su, Xiulong Liu, and Eli Shlizerman. Predict & cluster: Unsupervised skeleton based action recognition. In CVPR, pages 9631–9640, 2020.

[148] Shan Su, Jung Pyo Hong, Jianbo Shi, and Hyun Soo Park. Predicting behaviors of basketball players from first person videos. In CVPR, pages 1501–1510, 2017.

[149] Chen Sun, Abhinav Shrivastava, Carl Vondrick, Rahul Sukthankar, Kevin Murphy, and Cordelia Schmid. Relational action forecasting. In CVPR, pages 273–283, 2019.

[150] Dídac Surís, Ruoshi Liu, and Carl Vondrick. Learning the predictability of the future. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 12607–12617, 2021.

REFERENCES

[151] Tomoyuki Suzuki, Hirokatsu Kataoka, Yoshimitsu Aoki, and Yutaka Satoh. Anticipating traffic accidents with adaptive loss and large-scale incident db. In CVPR, pages 3521–3529, 2018.

[152] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions. In CVPR, pages 1–9, 2015.

[153] Yongyi Tang, Lin Ma, Wei Liu, and Weishi Zheng. Long-term human motion prediction by modeling motion context and enhancing motion dynamic. arXiv preprint arXiv:1805.02513, 2018.

[154] Du Tran, Lubomir Bourdev, Rob Fergus, Lorenzo Torresani, and Manohar Paluri. Learning spatiotemporal features with 3D convolutional networks. In ICCV, pages 4489–4497, 2015.

[155] Ioannis Tsochantaridis, Thorsten Joachims, Thomas Hofmann, and Yasemin Altun. Large margin methods for structured and interdependent output variables. Journal of Machine Learning Research, 6(9):1453–1484, 2005.

[156] Sergey Tulyakov, Ming-Yu Liu, Xiaodong Yang, and Jan Kautz. MocoGAN: Decomposing motion and content for video generation. In CVPR, pages 1526–1535, 2018.

[157] Ruben Villegas, Arkanath Pathak, Harini Kannan, Dumitru Erhan, Quoc V Le, and Honglak Lee. High fidelity video prediction with large stochastic recurrent neural networks. In NIPS, pages 81–91, 2019.

[158] Ruben Villegas, Jimei Yang, Seunghoon Hong, Xunyu Lin, and Honglak Lee. Decomposing motion and content for natural video sequence prediction. arXiv preprint arXiv:1706.08033, 2017.

[159] Ruben Villegas, Jimei Yang, Yuliang Zou, Sungryull Sohn, Xunyu Lin, and Honglak Lee. Learning to generate long-term future via hierarchical prediction. arXiv preprint arXiv:1704.05831, 2017.

[160] Carl Vondrick, Hamed Pirsiavash, and Antonio Torralba. Anticipating visual representations from unlabeled video. In CVPR, pages 98–106, 2016.

[161] Borui Wang, Ehsan Adeli, Hsu-kuang Chiu, De-An Huang, and Juan Carlos Niebles. Imitation learning for human pose prediction. In ICCV, pages 7124–7133, 2019.

[162] Dong Wang, Yuan Yuan, and Qi Wang. Early action prediction with generative adversarial networks. IEEE Access, pages 35795–35804, 2019.

[163] Heng Wang and Cordelia Schmid. Action recognition with improved trajectories. In ICCV, pages 3551–3558, 2013.

[164] Limin Wang, Yuanjun Xiong, Zhe Wang, and Yu Qiao. Towards good practices for very deep two-stream convnets. In ECCV, 2016.

[165] Xionghui Wang, Jian-Fang Hu, Jian-Huang Lai, Jianguo Zhang, and Wei-Shi Zheng. Progressive teacher-student learning for early action prediction. In CVPR, pages 3556–3565, 2019.

[166] Chao-Yuan Wu, Christoph Feichtenhofer, Haoqi Fan, Kaiming He, Philipp Krahenbuhl, and Ross Girshick. Long-term feature banks for detailed video understanding. In CVPR, pages 284–293, 2019.

REFERENCES

[167] Xinxiao Wu, Ruiqi Wang, Jingyi Hou, Hanxi Lin, and Jiebo Luo. Spatial–temporal relation reasoning for action prediction in videos. International Journal of Computer Vision, 129(5):1484–1505, 2021.

[168] Yue Wu, Rongrong Gao, Jaesik Park, and Qifeng Chen. Future video synthesis with object motion prediction. In CVPR, pages 5539–5548, 2020.

[169] Lu Xia, Chia-Chih Chen, and Jake K Aggarwal. View invariant human action recognition using histograms of 3D joints. In CVPRW, pages 20–27, 2012.

[170] Wanru Xu, Jian Yu, Zhenjiang Miao, Lili Wan, and Qiang Ji. Prediction-CGAN: Human action prediction with conditional generative adversarial networks. In MM, pages 611–619, 2019.

[171] Takuma Yagi, Karttikeya Mangalam, Ryo Yonetani, and Yoichi Sato. Future person localization in first-person videos. In CVPR, pages 7593–7602, 2018.

[172] Yufei Ye, Dhiraj Gandhi, Abhinav Gupta, and Shubham Tulsiani. Object-centric forward modeling for model predictive control. In Conference on Robot Learning, pages 100–109, 2020.

[173] Yufei Ye, Maneesh Singh, Abhinav Gupta, and Shubham Tulsiani. Compositional video prediction. In ICCV, pages 10353–10362, 2019.

[174] Serena Yeung, Olga Russakovsky, Ning Jin, Mykhaylo Andriluka, Greg Mori, and Li Fei-Fei. Every moment counts: Dense detailed labeling of actions in complex videos. IJCV, 126(2-4):375–389, 2018.

[175] Fisher Yu, Haofeng Chen, Xin Wang, Wenqi Xian, Yingying Chen, Fangchen Liu, Vashisht Madhavan, and Trevor Darrell. BDD100K: A diverse driving dataset for heterogeneous multitask learning. In CVPR, pages 2636–2645, 2020.

[176] Gang Yu, Zicheng Liu, and Junsong Yuan. Discriminative orderlet mining for real-time recognition of human-object interaction. In ACCV, pages 50–65, 2014.

[177] Tong Yu and Nicolas Padoy. Encode the unseen: Predictive video hashing for scalable mid-stream retrieval. arXiv preprint arXiv:2009.14661, 2020.

[178] Kuo-Hao Zeng, Shih-Han Chou, Fu-Hsiang Chan, Juan Carlos Niebles, and Min Sun. Agent-centric risk assessment: Accident anticipation and risky region localization. In CVPR, pages 2222–2230, 2017.

[179] Kuo-Hao Zeng, William B Shen, De-An Huang, Min Sun, and Juan Carlos Niebles. Visual forecasting by imitating dynamics in natural sequences. In ICCV, pages 2999–3008, 2017.

[180] Mengmi Zhang, Keng Teck Ma, Joo Hwee Lim, Qi Zhao, and Jiashi Feng. Deep future gaze: Gaze anticipation on egocentric videos using adversarial networks. In CVPR, pages 4372–4381, 2017.

[181] Zhengyou Zhang. Microsoft kinect sensor and its effect. IEEE multimedia, 19(2):4–10, 2012.

[182] He Zhao and Richard P Wildes. Spatiotemporal feature residual propagation for action prediction. In ICCV, pages 7003–7012, 2019.

[183] He Zhao and Richard P Wildes. On diverse asynchronous activity anticipation. In ECCV, pages 9925–9934, 2020.

[184] Bolei Zhou, Alex Andonian, Aude Oliva, and Antonio Torralba. Temporal relational reasoning in videos. In ECCV, pages 803–818, 2018.

[185] Xiaolong Zhu, Xuhui Jia, and Kwan-Yee K Wong. Pixel-level hand detection with shape-aware structured forests. In ACCV, pages 64–78, 2014.