

Generative models for learning and understanding dynamic scene activity

Hilary Buxton
School of Cognitive and Computing Sciences,
University of Sussex,
Brighton BN1 9QH,
UK.

Abstract— We are entering an era of more intelligent cognitive vision systems. Such systems can analyse activity in dynamic scenes to compute conceptual descriptions from motion trajectories of moving people and the objects they interact with. Here we review progress in the development of flexible, generative models that can explain visual input as a combination of hidden variables and can adapt to new types of input. Such models are particularly appropriate for the tasks posed by cognitive vision as they incorporate learning as well as having sufficient structure to represent a general class of problems. In addition, generative models explain all aspects of the input rather than attempting to ignore irrelevant sources of variation as in exemplar-based learning. Applications of these models in visual interaction for education, smart rooms and cars, as well as surveillance systems is also briefly reviewed.

I. INTRODUCTION

How can we develop *visual intelligence*? What are the basic research issues when we go from live video input to meaningful behaviour in our vision systems? We know that the gap cannot be closed by structural computer vision techniques alone. Such techniques are concerned with direct space-time aspects of the physical world and reconstruction of the observed scene. However, we are concerned with system purpose to integrate our vision modules into complete intelligent systems, which can understand observed activities as well as track the objects and people in the scene. Ballard's landmark paper [1] proposes that vision is best understood in the context of the visual behaviours engaging the system without requiring detailed internal representations of the scene. His approach also recognises that it is important to have a system framework that integrates visual processing within the task context. In general, it is not necessary to have perfect object reconstruction before we achieve the visual understanding required for real-world applications.

What do cognitive vision systems entail? The basic approaches combine techniques from symbolic or subsymbolic AI with computer vision techniques in some way. Naturally, we then encounter many of the major issues in AI such as knowledge representation and reasoning, control and the handling of uncertainty, as well as machine learning. Much of the work assumes that knowledge drives reasoning in visual interpretation (*seeing as*), thus visual context is seen as essential for understanding what is depicted

in images or image sequences. If we are to build efficient systems that can tackle many different tasks, high-level attention and control (*seeing for*) is also seen as essential. In addition, if we are to incorporate scene and task knowledge, we have to address the question of how such knowledge can be acquired. Knowledge structures were traditionally designed by hand but here we argue for the importance of automating their construction (*learning*) since learnt models are well suited to online analysis.

Learning in a vision system can be at the level of object models, their movements and actions, and how to control views and processing in the system. Some kinds of knowledge can be hand-crafted, for example, *explicit models* designed to provide a solution to one particular application. Such models are necessarily fixed after construction, with no flexibility to adapt to new kinds of data. A typical explicit model is given in early work by Hogg [2] on an articulated cylinder model of a walking man with spatio-temporal constraints on the movement patterns. This was advanced research for model-based vision in 1983, however, such models are slow to match to the image data and cannot adapt to changes in conditions in the online system. Even with more recent work where aspects of the model are learnt from training data, the fact that the model is explicit means that human intervention is required for labeling the examples.

An alternative approach has been developed more recently with learnt, *exemplar-based models* that have an associated distance metric. For this kind of model, it is necessary to remove irrelevant variation by preprocessing to derive the set of exemplars. A typical exemplar-based model is given in work by Howell and Buxton [3] on a set of prototypes from different views to support pose invariant identification. While such models are fast to match using the distance metric and can adapt, they are still crucially dependent on choice of preprocessing for the application.

Exciting new work on learning *flexible models* could take such work further, although so far only applied to more basic vision tasks such as segmentation, digit recognition, tracking and 'sprite' modeling [4]. Flexible models have evolved in the machine learning community and cover a wide class of parametric models for signal processing, which includes Principal Component Analysis (PCA) and Gaussian mixtures (VQ) as well as Hidden Markov Models (HMM) and Bayesian Belief Networks (BBN) etc. Here

TABLE I
GENERATIVE MODEL RELATIONSHIPS.

Initial	Extension	Final
Gaussian	mixture	VQ
Gaussian	reduce dimension	PCA
VQ	dynamics	HMM
PCA	independence	ICA
HMM	coupling	CHMM
HMM	variable length	VLMM
ICA	hierarchy	BBN
BBN	dynamics	DBN
DBN	utility	DDN

we are interested in those that are ‘generative’, that is, with probability distributions over hidden variables estimated from input data. A clear advantage of such models is that some knowledge of the complexity of the problem, such as number of dimensions of the hidden space, can be built in. These generative models can then adapt to the different types of input required by the application as well as predict and explain the data.

Essentially, what is proposed here is that a powerful way of tackling these issues of learning and inference as well as the use of context and active task control is to use generative models. In cognitive vision, there are two main classes of generative model which have been extensively applied: graphical models and deformable models. These are to a large extent complementary and can use well-established statistical learning theory in either offline or online learning. A unifying review of theory and techniques for such models from the machine learning perspective is given in Roweis and Ghahramani [5]. They explain the relationship between the models and even offer a generative model for generative models. Table I gives an overview of how the models can be regarded as extensions of each other.

Markov models and Bayesian networks are generally regarded as graphical models (for example Figure 1). The Gaussian is the underlying measurement model in automated learning of more structured generative models, where variability in the data can be explained (or predicted) from inferred (hidden) model parameters. HMM models can be regarded as extending Gaussian mixture (VQ) models by having learnt dynamic dependencies between states. These have a chain of simple dependencies on the immediately previous state. They can be extended to either 1) coupled dependencies with states in another HMM to form a Coupled Hidden Markov Model (CHMM) or 2) possible longer term temporal dependencies with previous states in the same HMM to form a Variable Length Markov Model (VLMM). Independent Component Analysis (ICA) models can be regarded as extending PCA models, commonly used to characterise data by a reduced set of dimensions, with the additional constraint of statistical independence of underlying model parameters. BBN models have this statistical independence property plus a possible hierarchy of dependencies between the hidden model vari-

ables. They can be extended by having learnt dynamic dependencies between states to form a Dynamic Bayesian Network (DBN) of some kind, e.g. with a simple Markov relationship over time. Adding utility theory allows decision support for rational agents using utility nodes and decision nodes for actions to form a Dynamic Decision Network (DDN) of some kind, although this extension is usually hand-coded.

II. GRAPHICAL MODELS

Probabilistic frameworks, in common with more general AI, have much to offer in dealing with the pervasive problem of uncertainty in visual evidence, and allow full information integration as proposed by Pearl [6]. The representation of constraints in a Bayesian Belief Network, BBN or ‘Bayes net’, can be achieved by mapping into the structure of a Directed Acyclic Graph so that the nodes represent concepts of interest and dependencies are given by causal links. Bayes nets can be learnt and model dependencies for either static (BBN) or dynamic (DBN) domains as well as incorporate decision theory in Dynamic Decision Networks (DDN). A simpler chained structure with single causal dependencies over time, the Hidden Markov Model (HMM), is often used for speech analysis [7] and has been extensively adapted for analysis of dynamic scenes and perceptual control as described below. In probabilistic reasoning [8], the likelihood of classes of objects or events is inferred by propagation of belief values in the light of changing evidence. Early work incorporating Bayes nets was developed by Levitt and Binford [9], [10] to make model-based vision reliable, while remaining computationally tractable.

Bayes nets have been widely adopted in vision systems as they are applicable to all levels of processing due to fast numerical updating in singly connected trees. There are techniques to decompose complex models and handle networks with multiple causes as well as learn the parameters for such networks [11]. Rimey and Brown introduced such techniques for active vision, with both control of camera movement and the selective processing required in task-based perceptual control [12], [13]. As discussed by Gong and Buxton [14], Bayes nets provide a clear way to map contextual constraints from the scene onto the computation of the visual interpretation by combining known causal dependencies with estimated statistical knowledge. They are essentially providing closed-loop control using both top-down and bottom-up messages in the propagation of belief values. They also provide the possibility of learning and refining visual representations by observation [15], [16]. Bayes nets have been used in many demanding applications such as BATmobile [17] and TEA system [13].

HMMs are also widely used in visual processing, as seen in the review of recent work on behaviour analysis below. The advantage here is that the ‘hidden’ purposes of regular behaviour patterns can be learnt from examples, i.e. the structure of the model as well as the parameters are easily learnt using the Baum-Welch (or EM) algorithm [7]. For more general Bayes nets, the dependency structure may be unknown, which complicates the learning process [18] but

it can be solved by using ‘structural’ EM with local search in the M step [19]. Conditional probability learning for the state transitions in HMMs is straightforward, for example, in early work by Gong [20], [21] the movement patterns of vehicles on an airport ground-plane were learnt and provided a generative model used in prediction. More recent work uses *coupled* CHMMs to learn models for interactions [22], *parameterised* HMMs for gesture interpretation [23], and *variable length* VLMMs for virtual reality systems [24].

A. Formalism

Bayesian belief networks are Directed Acyclic Graphs (DAG) in which each node represents an uncertain quantity using variables with multi-possible values. The arcs connecting the nodes signify the direct causal influences between the linked variables with the strengths of such influences quantified by associated conditional probabilities. If we assume a variable in the network is X_i , and a selection of variables Π_{X_i} are the direct causes of X_i , the strengths of these direct influences are quantified by assigning the variable X_i a link matrix $P(x_i|\Pi_{X_i})$, given any combination of instantiations of the parent set Π_{X_i} . The conjunction of all the local link matrices of variables X_i in the network (for $1 \leq i \leq n$ where n is the total number of the variables) specifies a complete and consistent global model which provides answers to all the probabilistic queries. Such a conjunction is given by the overall joint distribution function over the variables X_1, \dots, X_n : $P(x_1, x_2, \dots, x_n) = \prod_{i=1}^n P(x_i|\Pi_{X_i})$, where lower case symbols stand for a particular instantiation of the corresponding variables.

In a belief network, if we quantify the degree of coherence between the expectations (\mathbf{X}) and the evidence (\mathbf{e}) by a measure of local belief $BEL(x) = P(x|\mathbf{e})$, and define belief commitments as the tentative acceptance of a subset of hypotheses that together constitute a most satisfactory explanation of the evidence at hand, then, Bayesian belief revision amounts to the updating of belief commitments by distributed local message passing operations. Instead of associating a belief measure with each individual hypothesis locally, belief revision identifies a composite set of hypotheses that best explains the evidence. We call such a set the Most-Probable-Explanation (MPE). In computational terms, this means finding the most probable instantiations of all hypothetical variables given the observation.

Let \mathbf{W} stand for all the variables concerned, inclusive of those in \mathbf{e} , any particular instantiation of variables in \mathbf{W} that is also consistent with \mathbf{e} will be regarded as an *extension* or *explanation* of \mathbf{e} . The problem then is to find an extension \mathbf{w}^* that maximises the conditional probability $P(\mathbf{w}|\mathbf{e})$. In other words, $\mathbf{W} = \mathbf{w}^*$ is the MPE of the evidence if $P(\mathbf{w}^*|\mathbf{e}) = \max_{\mathbf{w}} P(\mathbf{w}|\mathbf{e})$. Here, \mathbf{w}^* is obtained by first locally computing the belief function for each variable X mentioned above, i.e. $BEL^*(x) = \max_{\mathbf{w}'_X} P(x, \mathbf{w}'_X|\mathbf{e})$ where $\mathbf{W}'_X = \mathbf{W} - X$ and second, propagating local messages. The local messages are defined as: if X has n parents U_1, U_2, \dots, U_n and m children Y_1, Y_2, \dots, Y_m , then node X receives messages $\pi_{X_i}^*(u_i), i = 1, \dots, n$ from its parents

and $\lambda_{Y_j}^*(x), j = 1, \dots, m$ from its children given by:

$\pi_{X_i}^*(u_i)$ is the probability of the most probable tail-extension of the hypothetical value $U_i = u_i$ relative to the link $U_i \rightarrow X$ and is known as an *explanation*,

$\lambda_{Y_j}^*(x)$ is the conditional probability of the most probable head-extension of the hypothetical value $X = x$ relative to the link $X \rightarrow Y_j$, known as a *forecast*.

It is important to understand the conceptual essence of this propagation mechanism. For each hypothetical value of a variable X , there exists a best extension of the complementary variables \mathbf{W}'_X . The problem of finding the best extension of $X = x$ can be decomposed into finding the best complementary extension to each of the neighbouring variables according to the conditional independence between X and the rest. This information can then be used to decide the best instantiation of X . This process of decomposition is applied recursively until it reaches the network boundary where evidence variables have predetermined values.

B. Tracking objects

Blake and Isard learn to track visual contours for general applications [25], without using a graphical model but with the possibility of building constraints into the prior term to influence interpretation. However, it can be argued that for application-specific tracking, there should be modeling of dependencies and explicit association of entities over time [26]. More recently, Black and Fleet [27] have developed a full generative Bayesian framework for tracking motion boundaries. Buxton and Gong [15] have developed a systematic methodology for the design and integration of advanced vision systems using Bayes nets. These networks allow dynamic updating of values in evidence and interpretation nodes, but not specification of the temporal constraints themselves. Howarth and Buxton, as discussed below, used dynamically reconfigured nets to model evolving spatial relationships of vehicles as they move through the scene before using a standard ‘tasknet’ BBN in the behaviour evaluation. The usual approach is to use map-like knowledge of the environment to develop expectations of likely object motion, as in Buxton and Gong [14], for the segmentation and tracking of vehicles. For example, in Figure 1 the top node of the graph is image grid position which causally affects two daughter nodes coding orientation and size of the vehicle. These, in turn causally affect the observed flow vector features to be grouped under these contextual constraints for tracking. Others [17] have adopted dynamic probabilistic networks [28], which make use of the simple Markov property that the future is independent of the past given the present state.

C. Tracking humans

Isard and Blake have continued to develop general approaches to tracking that are well-suited to tracking fast human movement where it is particularly important to keep multiple likely explanations of motion patterns [29], [30]. Conditional Expectation Maximisation, (CEM) or ‘condensation’, trackers of this kind have their roots in recursive filtering and control theory and go well with a probabilistic

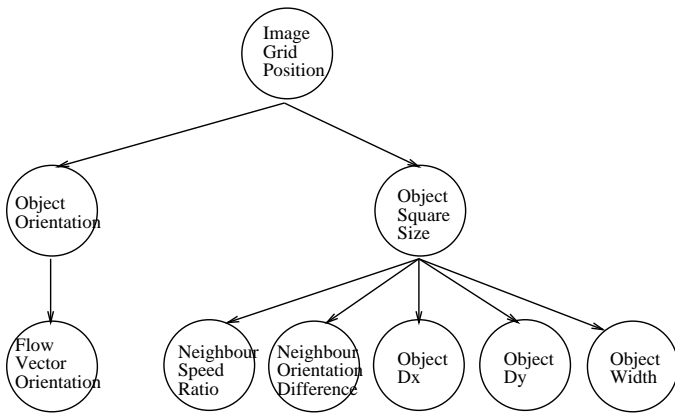


Fig. 1. Belief network that captures dependent relationships for motion segmentation and tracking (from [14]).

approach to behaviour analysis. They provide a framework for representing density as a set of samples drawn from the underlying distribution in a Bayesian manner. The main difficulty is that these trackers must maintain a large number of samples, e.g. 15,000 to track a hand while drawing, in order to fully represent the underlying density. Recent developments by Hogg and colleagues (see next section) have improved on this, but in general we should note that where the dynamics are known, they can be modeled to simplify tracking, for example Wren [31].

D. Behaviour analysis

Some of the recent work using graphical models in behaviour analysis has been mentioned above. For example, support of spatial and temporal reasoning to generate dynamic scene descriptions [32], which uses hand-coded Bayesian networks. Bayesian systems can also be used for intelligent vehicle control as in the BATmobile [17] or recent work by Oliver and Pentland [33]. We have also briefly introduced the coupled HMMs of Brand and colleagues [22] for understanding behaviour interactions, although this approach requires a great deal of training data. This is also true of parameterised HMMs [23], which can suffer from lack of stability in the interpretation compared to deformable model tracking and analysis seen in the next section. More recent work by Oliver, Rosario and Pentland [34], [35] has developed reliable Bayesian vision systems. Two exciting recent developments are: 1) work by Galata, Johnson and Hogg using deformable models with HMM behaviour models for virtual actors [24] and 2) the action reaction learning of Jebara and Pentland [36], [37], which models interactions and exploits new ideas from Support Vector Machines (SVMs) in conjunction with generative Bayesian theory.

D.1 Visual representation

The main system we look at in some detail here, as an exemplar of graphical models, is the HIVIS-watcher system [38], [39]. When reasoning about the behaviour of dynamic objects, it is useful if the representation of the properties related to each object are described in a local relative coordinate system. This involves recognising each

moving object so that an intrinsic-front can be identified together with its spatial extent for attentive processing. In a surveillance system we can obtain the poses of the scene objects via model-matching making local reasoning attractive. The *local-form* is representation and reasoning that uses the intrinsic frame-of-reference of a perceived object. The *global-form* is representation and reasoning that uses the perceiver’s frame-of-reference, which operates over the whole field-of-view. In active surveillance where we are inherently concerned with the “here-and-now”, it is important to form a consistent, task relevant interpretation of the observed behaviour. The system achieves this by task-level control policies that use typical-object behaviour models to specify both the preattentive (global) and attentive (local) processing to be performed in watching for evolving behaviour interactions. In particular, deictic spatio-temporal relationships are used to describe evolving contexts for observer and scene objects [40]. In this situated approach, it is not necessary to name and describe every object, but register only those relevant to the task so that the information registered is proportional to properties of interest rather than everything derivable in the dynamic scene.

D.2 Visual control

In HIVIS-watcher, there are three separate elements: the ‘virtual-world’ which holds data about the world, the ‘peripheral-system’ with operators that access the world, and the ‘central-system’ which controls system behaviour. The peripheral-system is based on Ullman’s [41] visual routine processor, with reactive planning following the approach of Agre and Chapman [42] using spatial indexes [43]. Event detection operators are only run when selected by the task-level control system in this kind of situated vision approach [44]. The operators in the peripheral-system are separated into preattentive ones that are global, simple, and of low-cost and attentive ones which are applied to a single object and are more complex. The preattentive operators are used to guide application of attentive ones. Example preattentive operators include *gross-change-in-motion* and *mutual-proximity*, Howarth and Buxton [45], [38]. The motivation behind the preattentive and attentional cues chosen here, is their usefulness on low-level data such as the identification of possible objects from clustering flow-vectors (see Gong and Buxton [46]). If these movement descriptions comply with the preattentive cue, then they become candidates for further attentional processing, such as model-matching, to obtain deictic relationships about the objects. Basically, this strategy for perceptual control simply finds where something interesting is, then works out what is happening in detail.

Bayes nets are the main mechanism used to update the context for interpretation of the dynamic scene at the behaviour level. This framework allows knowledge representation (symbolic and probabilistic) as well as playing a role in the online control of visual processing. For example, *gross-change-in-motion* is the preattentive cue to watch for in analysis of dynamic ‘giveaway’ behaviour and to start gathering detailed evidence for the Bayesian tasknet. Simi-

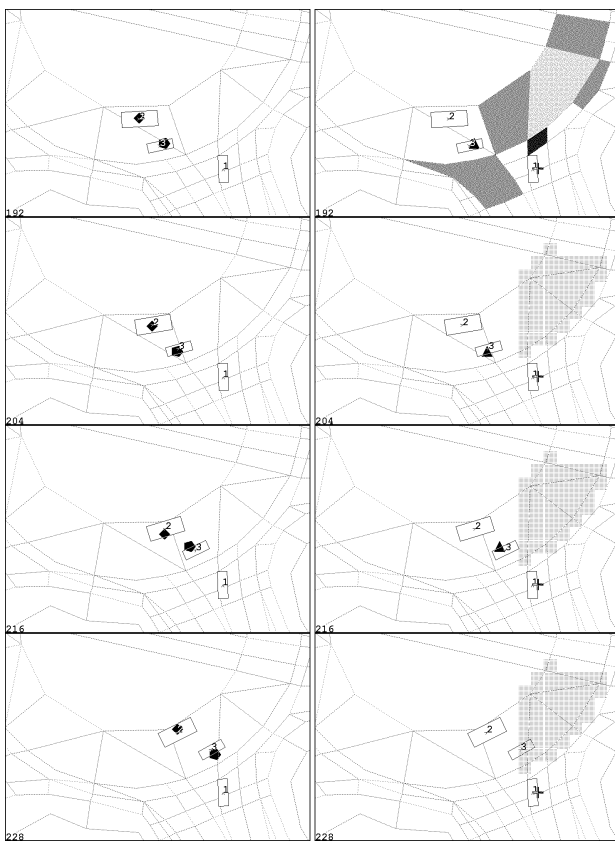


Fig. 2. Results from HIVIS-watcher (from [38]) showing overtaking policy (left) and giveaway policy (right).

larly, *mutual-proximity* is the cue for ‘overtaking’ or ‘following’ so these are treated as starting the attentive processing of evidence gathering when these are the active policy. Evidence in this system is in terms of deictic relationships such as “ref-obj3 is moving faster than ref-obj2”, which help confirm or reject decisions about the relevant behaviour. Prior probabilities are not very relevant in such decision making but statistical knowledge is learnt as conditional probabilities for the dynamic relationships, also how likely different classes are *a priori* to initialise interpretation is relevant at the movement level [14]. All Bayes nets are updated using both parent to child (expectation) and child to parent (evidence) to give maximum likelihood explanation *a posteriori*. The complete dynamic updating cycle plays a role in controlling evidence collection as nodes can be parameters that affect the perceptual processing. For example, in Figure 2 the same data is processed in different ways depending on the control policy coded in a DDN.

D.3 Visual learning

The learning above is offline but is exploited online in the HIVIS system. So far, the processing is structured by analysis of what was involved in verifying that some behaviour is taking place, both what should be attended (via the preattentive cues and attentional markers) and how the reasoning should be structured using *typical object behaviour models*. However, this hand-coding is very difficult and involves a lot of empirical evaluation in order

to get an effective processing scheme. Ongoing work learns the online visual cues for behaviours using a mixture of connectionist and Bayesian learning techniques to develop open classes of behavioural models. For complex models it helps to learn the detailed weights using Probabilistic Neural Networks (PNNs) from an initial DAG structure. New work by Frey and Jojic [4] on flexible models as well as learning with variational methods in the graphical models community [47], [48] is rapidly developing techniques which can be applied to learn structure as well as parameters for cognitive vision tasks. For example, Yacoub and Black [49] have developed modeling and recognition with PCA ‘activity-bases’.

III. DEFORMABLE MODELS

Original work by Terzopoulos [50] pioneered the use of deformable models in realistic simulations of human heads using biomechanical facial muscles and skin movements. This work led to early methods for shape and non-rigid motion estimation going from synthesis to an interpretive task [51]. An alternative approach was developed by Edwards and colleagues, based on work by Cootes and Taylor [52]. This simply models the appearance of the face in 2D with parameterised texture variations and 2D deformations to simulate different expressions and viewpoints [53] rather than model the underlying 3D structure. Such models can provide good graphical representations for synthesis of appearance and motion but reversing this for interpretation is less well understood. However, the Leeds group has shown we can both learn adaptive shape models for humans and use them in learning likely spatio-temporal patterns of activity for tracking, as in the early work by Baumberg and Hogg [54], [55]

So what is a deformable model and how is it used in this visual interpretation work? Physically-based vibration modes are fundamental for the description of such models and building in such known constraints into the perception process can allow difficult problems in computer vision to be solved. The constraints are mainly learnt from observation for visual interpretation, e.g. Baumberg and Hogg [54]. However, in general, these constraints need not even be based on real physical properties but just act to condition the interpretation. For tracking a non-rigid deformable object, early work by Pentland and Horowitz [56] showed how to recover motion and structure using ‘finite elements’. This method relies on assumptions about elasticity and density distribution with vibration modes derived via mass and stiffness matrices in the main equation governing the physical behaviour. Further work by Nastar [57] used such ‘modal analysis’ in a wider range of applications. However, it is the combination of acquisition of these modes of variation by observation and their use in generative perceptual control that has proved crucial for more general application in vision.

The use of training in deformable model analysis has much in common with training in neural networks. Cootes and Taylor [58], in their early work, introduced the Point Distribution Model (PDM) which is derived from a set of

characteristic training data for the problem at hand and parameterised by a set of orthogonal ‘modes of variation’. The training shapes are then represented by a subset of vectors which account for the majority of the observed variations. The PDM has proved extremely useful for image sequence analysis for tracking contours [54] and locating structures in medical images [59]. However, there was originally no intrinsic time dimension and, in early work, characteristic points for the training sets were selected by hand. More recent work by Baumberg and Hogg [55] successfully tackled both these drawbacks and has led to the development of not only tracking of walking people but also generalisation of the techniques using Probability Density Functions’ (PDFs) on extracted trajectories. These are used in event analysis (e.g. Johnson and Hogg [60]) and the full visual interaction modeling we describe below, which uses extended statistical learning techniques.

A. Formalism

Point distribution models are a well-established tool for statistical analysis of visual data. The PDM is based on a set of example shapes of a given object (or behaviour pattern if temporal extension). Traditionally each shape is defined by landmark points that are selected as important, corresponding to features of the object (or trajectory). This approach then allows a class of objects (or behavior patterns) to be characterised by a small set of parameters. The original linear PDM [58] was trained on a set of shapes with n landmark points aligned to the mean shape $\bar{\mathbf{x}}$. The distance from the mean is calculated using PCA to get the vector $\mathbf{dx} = \mathbf{x} - \bar{\mathbf{x}}$ for the n 2D landmark vectors $\mathbf{x} = (x_1, y_1, \dots, x_n, y_n)$. Next the $2n \times 2n$ covariance matrix $C = E(\mathbf{dx}\mathbf{dx}^T)$ is found using the expectation value E over the training examples. Eigenvectors of C then correspond to the variation modes for this data with the largest eigenvalue describing the most significant mode etc. The PDM model is simply the mean shape \mathbf{x} , together with the t eigenvectors $\mathbf{p}_1, \dots, \mathbf{p}_t$ characterising a sufficient number of significant modes of variation to describe the training data for the task at hand.

The PDM can be extended to parameterised curves such as the cubic B-spline [60]. They show that the point measurements can be mapped to the modes of variation using $\mathbf{b} = P^T(\mathbf{x} - \bar{\mathbf{x}})$, where \mathbf{b} is a vector of t shape parameters, \mathbf{x} is a vector of $2n$ positions and P is a matrix with columns corresponding to eigenvectors of the covariance matrix. Using a set of parameterised shapes for training the model, represented by n control points, allows the automated placing of a set of measurements at regular intervals along a continuous curve. However, in contrast to landmark points set by hand, any particular measurement will affect several control points, which violates the independence assumption implicit in the measurement model (Gaussian sampling noise with variance r). The formal solution to this problem is to use the covariance matrix $R = r\mathcal{H}^{-1}$, where $\mathcal{H}_{ij} = \int H_i(u)H_j(u)du$ and H_i is the interpolation function for the parameter (see [61]).

In practice a covariance matrix S is estimated from train-

ing data to obtain $2n$ eigenshapes using the relationship $S\mathcal{H}\mathbf{e}_i = \lambda_i\mathbf{e}_i$. The eigenvectors \mathbf{e}_i are made orthonormal to the inner product $\langle \mathbf{u}, \mathbf{u} \rangle = \mathbf{u}^T\mathcal{H}\mathbf{u}$. The shape vectors are defined by $\mathbf{u} = \sum_{i=1}^{2n} \mu_i \mathbf{e}_i + \bar{\mathbf{u}}$, with coefficient $\mu_i = \langle \mathbf{u} - \bar{\mathbf{u}}, \mathbf{e}_i \rangle$. Over the training set these coefficients can then be shown to be linearly independent and allow analysis similar to that in the traditional PDM. That is, the coefficients for each eigenshape mode are updated independently and a subset t can be used sufficient for the task. The model is equivalent to a dynamic system using Finite Elements with stiffness matrix given by S^{-1} and mass matrix given by \mathcal{H} .

B. Tracking objects

Deformable shape models, such as the PDM, are derived from a set of training data containing representative objects. The boundary of each segmented object is traced automatically and fitted with a cubic B-spline contour with consistent control points or fitted with an active ‘snake’ [62]. As we have seen, using the PDM approach above requires PCA on the spline data so that a small set of the largest components characterise the data. This model can be reparameterised to make it even more compact for the on-line tracking. Kalman filters can then be used to track objects using position, scale and orientation of the model in addition to the compact shape parameters. The iterated Kalman filter is used by Baumberg [63] to adaptively set these parameters for rapid convergence in the tracking. Using the ground-plane constraint allows the 2D height and position parameters to be mapped into approximate world coordinates. This approach was developed for deformable car models by Sullivan and colleagues [64] and brought together for both humans and cars in a joint project [65]. In the combined system, cars can stop moving and be incorporated into the background by using a median filter over time, thus allowing the detection of moving objects in front of these stationary vehicles.

C. Tracking humans

A promising scheme using using PDMs has been mentioned above and tested for tracking human figures in image sequences, starting with the early work of Baumberg and Hogg [55], [66]. The flexible 2D prior model of the human outline shape can be used to recognise and track pedestrians in dynamic scenes. This uses a completely automated learnt model then in contrast to the explicit, hand-built 3D articulated model used in human tracking by Hogg in 1983 [2]. The Leeds deformable shape model incorporates the observed modes of variation as constraints so that tracking can be performed irrespective of partial occlusions and image noise. This model captures the apparent changes in shape found in images due to underlying human pose and deformation variation, as well as relative viewpoint variation of the walking human with respect to the camera in the 3D scene. Wren and Pentland have also built deformable models with a region-based approach in which moving ‘blobs’ are tracked in real-time, used in their ‘Pfinder’ system [67]. Further work by Wren allows more

detailed analysis of human body movements which incorporates some real physical constraints due to the skeletal structure [31]. This is closer to the work using biomimetic models of Terzopoulos and colleagues cited above, which incorporates models of the underlying muscle structure and elasticity.

D. Behaviour analysis

The main systems here to illustrate the deformable models approach are taken from the work of Johnson and Hogg [60], Morris and Hogg [68] and Galata, Johnson and Hogg [24]. Their work emphasises acquisition and encoding of spatial, temporal and procedural knowledge by passive observation of video sequences of typical interactions. The basis for the work is development of spatio-temporal models automatically learnt from video training data. These models have no explicit or hand-crafted element and are usually low dimensional linear transformations of the image space. These visual models of interaction have been applied to both surveillance applications and human computer interaction.

A basic component of these systems is representation of a given target activity by training a model with the video sequences depicting repeated examples of this target activity. The sequences must contain examples that span the range of ways in which the activity may be carried out and there will typically need to be hundreds of these examples. If the activity is on the ground-plane, such as pedestrians crossing a pathway or carpark, the constructed models must be sufficiently detailed to evaluate whether the pedestrian is deviating from the model and to allow prediction of the immediately following trajectory, up to some choice-point or junction for example. This kind of model, then, would be of use in surveillance systems to raise alarms and provide short-term occlusion handling in tracking.

D.1 Visual tracking

The *visual tracker* of Baumberg and Hogg [55], [66] is used in analysis of this kind and closely related to the active shape models of Cootes and Taylor [52] and active contours of Blake and Isard [25]. A closed B-spline contour is used to represent the image profile and consists of a set of control points. There is initial segmentation using image differencing to extract the moving objects for training, then the B-spline contour is wrapped around each extracted profile so that the first control point is located in a similar position in each example for training. The number and spacing of the control points is also held constant so that variations in control point location from example to example are due to observable, intrinsic shape variations. An eigenshape model then consists of the largest principal components that account for these variations. For example, about 10-20 components were sufficient for the walking pedestrian, with the most salient feature being the gap appearing and disappearing between the legs.

Tracking using this kind of model can provide a set of example trajectories by using a Kalman filter to constrain



Fig. 3. Results from tracking using deformable model to show trajectory extracted for further analysis (from [60]).

the search from frame to frame. For example, in Figure 3 a trajectory from such tracking is shown ready for further processing. In particular, these trajectories are then modeled using statistical learning techniques, this time to develop a full spatio-temporal model. This is achieved in two stages, first, in terms of instantaneous locations and velocities in the image plane using a 4D state vector. The set of state prototypes provides a simple behavior model capable of very short term prediction, dependent on immediately previous state rather than any history. The second stage then builds the extended behaviours such that associated activation levels increase and decrease along the evolving trajectory, held in a behaviour vector. The collection of behaviour vectors is then modeled as a distribution of a set of prototypes, derived by vector quantisation. This final set of prototypes is a compressed model of the observed behavior obtained by increasing abstraction of the data. For more details see Johnson and Hogg [60], who demonstrate its use in assessing the typicality of a given trajectory from the PDF.

D.2 Visual interaction

The *visual interaction* of Morris and Hogg [68] builds on this work to allow the capability of recognising, for example, ‘suspicious’ behaviour. Trajectory landmarks are found using closest approach to the nearest object and the distribution of speed and distance with respect to these landmarks, called an ‘event’, is used to characterise the underlying typical behaviour during the training phase. In the recognition phase, the probability of a given event is assessed using simple statistical tests based on observing the measured speed and distance, where low speed and distance are associated with suspicious behaviour. That is, moving quickly down a row of cars is not unusual nor is standing still a long way off but being both close and near stationary is noted as atypical (low probability). The trajectories are first observed as an ordered sequence of events but are resorted in order of increasing probability for assessment. The last stage requires characterisation of these increasing probabilities under the expectation that people get out of their cars and walk to the exit. This will start with low probability events when getting out of the car but rapidly increase if all other interactions are at a rea-

sonable speed and distance. However, if atypical behaviour is taking place, there will be a large number of low probability events. Supervised learning is used to provide the classification boundary.

D.3 Virtual reality

The *virtual reality* systems of Johnson, Galata and Hogg [69] also build on this approach to model interaction between people for applications in human computer interaction. Here a probabilistic model of the joint behaviours is learnt by observing sets of typical interactions offline in the training phase much as above but using the ‘eigen-shape’ B-spline profile as the base [66]. The contours of the two interacting people are then represented by the concatenated set of control points from the B-spline (third order) contours. This time the state vector uses separation and height of the left and right individuals together with the first derivative of these measures. From the data, a set of prototype states are again derived by vector quantisation and linked into a Markov chain to form the behaviour model. As in the surveillance work, the model can be used in a variety of ways in the on-line recognition and interpretation phase. For example, extrapolation forward in time from a tracked behaviour is obtained probabilistically from the model in order to generate a set of more or less likely continuations. The model can also fill in missing parts of the behaviour interaction using a Bayesian framework in a similar manner to representing and updating the *a posteriori* density used in the tracking algorithm [61]. In general, the set of plausible state hypotheses can be found from this density function, where the maximum represents the most likely hypothesis for the state of the interaction.

The main application of their work is in the synthesis of a virtual partner and this has been demonstrated for handshaking and turn-taking in speaker interaction. Further work in this direction by Galata, Johnson and Hogg [24] has developed this approach with the deformable model combined with a graphical model, a variable length Markov model. This can be applied to highly structured behaviour such as dance, aerobics and sign language. The process involves first segmenting the video stream into plausible behaviour segments and then using the variable length Markov model to represent the relationships between these atomic segments. The acquisition is again automated learning in the training phase and synthesis of behaviour for a virtual actor is the main application of the learnt models. As we saw in the section on graphical models, there are many variants of using an HMM and this approach then is close to work coupling dynamic models with the Markov chain representing long-term constraints [70], [71]. However, the extension here uses VLMMs to overcome problems in the iterative optimisation found in learning standard models where there can be local maxima. The flexible models first encode sequences of about 20ms and then the higher level sequences of atomic behaviours of about 1 second.

IV. APPLICATIONS

Generative models can be used to advance cognitive vision for applications in human-computer interaction, education and tutoring, smart rooms and cars, and surveillance systems.

Many researchers are developing useful techniques for the rapidly growing area of multimodal and multimedia interaction. Turk’s work at Microsoft with visual control of virtual actors is one example [72], as is Blake’s use of sophisticated multi-object trackers [73] in hand tracking [74] and new work by Frey and Jojic on learning flexible sprites [4]. Another well-known centre for such work is the MIT Media lab led by Pentland. Some of this research has been mentioned in the reviews earlier, especially under the graphical models section. MIT Media lab is also very active in the development of education and tutoring systems. For example, there has been ongoing research on sign language interpretation by Starner and Pentland [75] using HMMs. Bobick and others developed the exciting KidsRoom project while at MIT [76]. This was a perceptually-based environment in which children could interact while playing in a story-telling scenario. The EU Puppet project also uses story-telling in a virtual theatre for education as in Scaife and Rogers [77], where generative design of a virtual environment supports learning in children.

Smart rooms and intelligent environments are now a rapidly growing area with workshop series devoted to this topic. A visionary overview of general issues and some early research is given by Pentland [78]. A good example is given by the intelligent studios of Pinhanez and Bobick [79] who use their approximate world models and incorporate linguistic information to select vision routines and control of camera views. Smart cars using new sensor technology for vehicle control are being developed in conjunction with traffic monitoring in intelligent highway systems by Malik and colleagues [80], [81]. Also, a SmartCar testbed with learning techniques to model and recognise driving behaviour has been developed by Oliver and Pentland [33]. The techniques here are again exploiting graph-based models, both HMM and CHMMs for: *passing, changing lanes, turning, starting and stopping* with prediction a full second before the manoeuvre starts. This anticipation is essential for control and driver assistance systems.

A great deal of cognitive vision research has been done in the field of surveillance, for example early work on the EU VIEWS project discussed in the graphical models section [82]. Here again there is the need to integrate vision and language to deliver conceptual descriptions of the scene, which means that language must be grounded in vision and there must be a clear ontology for events, activities and scenarios. For example, recent work by Nevatia’s group has used a Bayesian approach to develop multi-agent event recognition [83], [84]. In addition, Bayes nets and dynamic decision nets have been developed for task-based control [85] and many of the reasoning tasks required in advanced surveillance systems by Buxton, Gong and Howarth [15], [86], [39]. These allow ongoing interpretation of the visual evidence in a control loop to actively direct processing for

tracking or behaviour evaluation and deliver a conceptual description of the dynamic scene which is relevant to the surveillance task.

V. CONCLUSION

To summarise the arguments in the earlier sections, we are suggesting that generative models are essential in the design and integration of cognitive vision systems. For the interpretation of a dynamic scene, these systems not only track and analyse movements of objects and people in the scene, but also evaluate their behaviour in the context of the vision system task. Thus, models are used in both interpretation and control of the visual processing. In addition, these models should ideally be learnt for the interpretation and control of cognitive vision tasks to automate construction and support adaptive processing.

A useful approach to cognitive vision system design with full information integration is offered by graphical models, where the representations reflect dynamic dependencies for interpretation of visual evidence. We reviewed work where Bayes nets and HMMs are partially structured using contextual knowledge [14] and detailed parameters learnt for prior and conditional probabilities between the important concepts in dynamic scene interpretation. Complex models suffer from the combinatorics of top-down and bottom-up message passing in online evaluation of beliefs required to compute the most likely explanation of the visual evidence. However, if singly connected trees of limited depth are used in a decomposition or simplification of the problem, they are effective for real-time systems, incorporating active selection for behaviour interpretation [38]. The contextual constraints and control of processing in updating beliefs from evidence in these graphical models is easily understood. Modeling essential aspects of situated vision by using deictic markers in conjunction with the Bayes nets and task-based control [39] is also possible. This type of approach has all the required characteristics for cognitive vision: contextual processing and control with learning capabilities.

We also reviewed some sophisticated work using the deformable models approach. Flexible, adaptive learning used statistical techniques to extract the models of typical behaviour at the level of movement trajectories, events, and more abstract sequences of movement. However, note that all this analysis was in the image plane and does not involve active visual processing or task control. There are approaches using fuller biomechanical models by Terzopoulos and colleagues [87] aimed more at an artificial life level of modeling for behaviour. These biomimetic agents can be situated in realistic environments to develop a software approach to the design of active vision systems [88]. Nevertheless, integration of cognitive vision systems involves more than the deformable models as they must be combined with some kind of situated vision framework to allow action or within a full cognitive vision framework to allow long-term reasoning and evaluation of behaviour. Thus, there are difficulties with both flexible, task-based control of processing and with the full use of scene context using de-

formable models. This means that these models are very appropriate for interpretation up to the level of movement but they have limitations as an approach to the design and integration of cognitive vision systems.

Finally, for applications, it seems that there is a need for interdisciplinary work in cognitive science, HCI and AI approaches to vision. For example, Thibadeau used animated sequences of simple geometrical shapes to demonstrate that human subjects offer intentional descriptions of the observed movement patterns [89]. With real people in real scenarios, human observers are able to predict criminal acts from video information leading up to the act [90]. The results here suggest that not only can correct predictions can be made, but also that it's possible to identify predictive behaviours. However, it remains for future cognitive studies using modern motion capture and animation to systematically explore the precise movement cues used in prediction.

We have seen how cognitive science research on deictic descriptions and attentional markers in situated vision [42], [43] has been used in the work of Howarth and Buxton [40]). This is particularly important for the generation of conceptual descriptions of dynamic scenes [39]. We also mentioned the role of active viewpoint control in cognitive vision systems, which is an important area of research but not reviewed here (see [91]). Recent research on overt attention in natural tasks, such as Land and McLeod's study of a cricket batsman trying to hit the ball [92], gives great insight into perceptual performance for such active camera systems. Future progress in this area depends on multidisciplinary projects that combine such cognitive studies with learnt generative models to develop effective cognitive vision systems.

ACKNOWLEDGEMENTS

The author gratefully acknowledges the invaluable discussion and help of co-workers Shaogang Gong and Richard Howarth who completed much of the original work in the section on graphical models reviewed here. Also thanks to David Hogg for permission reproduce figures and discussion of approaches for the section on deformable models. Finally, thanks are due for funding under EU IST cognitive vision initiative ActIPret project.

REFERENCES

- [1] D. Ballard, "Animate vision," *Artificial Intelligence*, vol. 48, pp. 57-86, 1991.
- [2] D.C. Hogg, "Model-based vision: A program to see a walking person," *Image and Vision Computing*, vol. 1, pp. 5-20, 1983.
- [3] A.J. Howell and H. Buxton, "Learning identity with Radial Basis Function networks," *Neurocomputing*, vol. 20, pp. 15-34, 1998.
- [4] B.J. Frey and N. Jovic, "Flexible models: A powerful alternative to exemplars and explicit models," in *CVPR Workshop on Models versus Exemplars*, Kauai, Hawaii, 2001.
- [5] S. Roweis and Z. Ghahramani, "A unifying review of linear Gaussian models," *Neural Computation*, vol. 11, pp. 305-345, 1999.
- [6] J. Pearl, *Probabilistic Reasoning in Intelligent Systems, Networks of Plausible Inference*, Morgan Kaufmann, 1988.
- [7] L.R. Rabiner, "A tutorial on hidden Markov models," *Proceedings of the IEEE*, vol. 77, pp. 257-286, 1989.
- [8] E. Charniak, "Bayesian networks without tears," *AI Magazine*, vol. 12, no. 4, pp. 50-63, 1991.

- [9] T.S. Levitt, J.M. Agosta, and T.O. Binford, "Model-based influence diagrams for machine vision," in *Uncertainty in Artificial Intelligence*. Machine Intelligence and Pattern Recognition Series, Volume 10, North-Holland, 1990.
- [10] T.S. Levitt, T.O. Binford, and G.J. Ettinger, "Utility-based control for computer vision," in *Uncertainty in Artificial Intelligence*. Machine Intelligence and Pattern Recognition Series, Volume 9, North-Holland, 1990.
- [11] D.J. Spiegelhalter and R.G. Cowell, "Learning in probabilistic expert systems," in *Bayesian Statistics 4*. Oxford University Press, 1992.
- [12] R.D. Rimey and C.M. Brown, "Where to look next using a Bayes net: Incorporating geometric relations," in *European Conference on Computer Vision*, Genoa, Italy, 1992, pp. 542-550.
- [13] R.D. Rimey and C.M. Brown, "Control of selective perception using Bayes nets and decision theory," *International Journal of Computer Vision*, vol. 12, pp. 173-209, 1994.
- [14] S. Gong and H. Buxton, "Bayesian nets for mapping contextual knowledge to computational constraints in motion segmentation and tracking," in *British Machine Vision Conference*, Guildford, UK, 1993, pp. 229-238.
- [15] H. Buxton and S. Gong, "Visual surveillance in a dynamic and uncertain world," *Artificial Intelligence*, vol. 78, pp. 431-459, 1995.
- [16] S.D. Whitehead and D.H. Ballard, "Learning to perceive and act by trial and error," *Machine Learning*, vol. 7, pp. 45-83, 1991.
- [17] J. Forbes, T. Huang, K. Kanazawa, and S. Russell, "The BAT-mobile: Towards a Bayesian automated taxi," in *International Joint Conference on Artificial Intelligence*, Montreal, Canada, 1995, pp. 1878-1885.
- [18] D. Heckerman, "A tutorial on learning with Bayesian networks," in *Learning in Graphical Models*, pp. 301-354. NATO Science Series, 1998.
- [19] N. Friedman, "Learning belief networks in the presence of missing values and hidden variables," in *International Conference on Machine Learning*, 1997, pp. 125-133.
- [20] S. Gong, "Visual behaviour: Modelling "hidden" purposes in motion," in *International Conference on Neural and Stochastic Methods in Image and Signal Processing*, San Diego, CA, 1992, pp. 45-57.
- [21] S. Gong, "Visual observation as reactive learning," in *International Conference on Adaptive and Learning Systems*, Orlando, FL, 1992, pp. 175-187.
- [22] M. Brand, N. Oliver, and A. Pentland, "Coupled hidden Markov models for complex action recognition," in *IEEE Conference on Computer Vision and Pattern Recognition*, San Juan, Puerto Rico, 1997.
- [23] A. Bobick and A. Wilson, "A state-based technique for the summarization and recognition of gesture," in *IEEE International Conference on Computer Vision*, Cambridge, MA, 1995, pp. 382-388.
- [24] A. Galata, N. Johnson, and D.C. Hogg, "Learning variable length Markov models of behaviour," *Computer Vision and Image Understanding*, vol. 81, pp. 398-413, 2001.
- [25] A. Blake, M. Isard, and D. Reynard, "Learning to track the visual motion of contours," *Artificial Intelligence*, vol. 78, pp. 179-212, 1995.
- [26] A.E. Nicholson and J.M. Brady, "The data association problem when monitoring robot vehicles using dynamic belief networks," in *European Conference on Artificial Intelligence*, Vienna, Austria, 1992, pp. 689-693.
- [27] M.J. Black and D.J. Fleet, "Probabilistic detection and tracking of motion boundaries," *International Journal of Computer Vision*, vol. 38, pp. 231-245, 2000.
- [28] T. Dean and K. Kanazawa, "Probabilistic temporal reasoning," in *American Association for Artificial Intelligence*, St Paul, Minnesota, 1988, pp. 524-528.
- [29] M. Isard and A. Blake, "Contour tracking by stochastic propagation of conditional density," in *European Conference on Computer Vision*, Cambridge, UK, 1996, pp. 343-356.
- [30] M. Isard and A. Blake, "A mixed-state condensation tracker with automatic model-switching," in *IEEE International Conference on Computer Vision*, Bombay, India, 1998, pp. 107-112.
- [31] C. Wren and A. Pentland, "Dynamic models of human motion," in *IEEE International Conference on Automatic Face and Gesture Recognition*, Nara, Japan, 1998.
- [32] H. Buxton and R. Howarth, "Spatial and temporal reasoning in the generation of dynamic scene descriptions," in *Representation and Processing of Spatial Expressions*, P. Olivier and K-P. Gapp, Eds. Lawrence Erlbaum Associates, 1997.
- [33] N. Oliver and A. Pentland, "Graphical models for driver behaviour recognition in a smartcar," in *Intelligent Vehicles*, Detroit, US, 2000.
- [34] N. Oliver, B. Rosario, and A. Pentland, "Graphical models for recognising human interactions," in *Advances in Neural Information Processing Systems*, Denver, Colorado, 1998.
- [35] N. Oliver, B. Rosario, and A. Pentland, "A Bayesian computer vision system for modeling human interactions," in *International Conference on Vision Systems*, Gran Canaria, Spain, 1999.
- [36] T. Jebara and A. Pentland, "Action reaction learning: Automatic visual analysis and synthesis of interactive behaviour," in *International conference on Vision Systems*, Gran Canaria, Spain, 1999, pp. 273-292.
- [37] T. Jebara and A. Pentland, "On reversing Jensen's Inequality," in *Advances in Neural Information Processing Systems*, Denver, Colorado, 2000, vol. 13.
- [38] R. Howarth and H. Buxton, "Visual surveillance monitoring and watching," in *European Conference on Computer Vision*, Cambridge, UK, 1996, pp. 321-334.
- [39] R. Howarth and H. Buxton, "Conceptual descriptions from monitoring and watching image sequences," *Image and Vision Computing*, vol. 18, pp. 105-135, 2000.
- [40] H. Buxton and R. Howarth, "Watching behaviour: The role of context and learning," in *International Conference on Image Processing*, Lausanne, Switzerland, 1996, pp. II:797-800.
- [41] S. Ullman, "Visual routines," *Cognition*, vol. 18, pp. 97-159, 1984.
- [42] P.E. Agre and D. Chapman, "Pengi: An implementation of a theory of activity," in *American Association for Artificial Intelligence*, 1987, pp. 268-272.
- [43] Z. Pylyshyn, "The role of location indexes in spatial perception: A sketch of the FINST spatial-index model," *Cognition*, vol. 32, pp. 65-97, 1989.
- [44] Z. Pylyshyn, "Visual indexes, preconceptual objects, and situated vision," *Cognition*, vol. 80, pp. 127-158, 2001.
- [45] R. Howarth and H. Buxton, "Selective attention in dynamic vision," in *International Joint Conference on Artificial Intelligence*, Chambéry, France, 1993, pp. 1579-1585.
- [46] S. Gong and H. Buxton, "On the visual expectations of moving objects: A probabilistic approach with Augmented Hidden Markov Model," in *European Conference on Artificial Intelligence*, Vienna, Austria, 1992, pp. 781-785.
- [47] M.I. Jordan, *Learning in Graphical Models*, NATO Science Series, 1998.
- [48] K. Murphy, "Learning Bayes net structure from sparse data sets," Tech. Rep., University of Berkeley, CA, 2001, Computer Science Technical Report.
- [49] Y. Yacoob and M.J. Black, "Parameterized modelling and recognition of activities," *Computer Vision and Image Understanding*, vol. 73, pp. 232-247, 1999.
- [50] D. Terzopoulos, "Physically-based facial modeling, analysis, and animation," *Journal of Visualisation and Computer Animation*, vol. 1, pp. 73-80, 1990.
- [51] D. Metaxas and D. Terzopoulos, "Shape and non-rigid motion estimation through physics-based synthesis," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 15, pp. 580-591, 1993.
- [52] T.F. Cootes, D. Cooper, C.J. Taylor, and J. Graham, "Active shape models - their training and application," *Computer Vision and Image Understanding*, vol. 61, pp. 38-59, 1995.
- [53] G. J. Edwards, C.J. Taylor, and T.F. Cootes, "Learning to identify and track faces in image sequences," in *British Machine Vision Conference*, Colchester, UK, 1997, pp. 130-139.
- [54] A. Baumberg and D.C. Hogg, "Learning flexible models from image sequences," in *European Conference on Computer Vision*, Stockholm, Sweden, 1994, pp. 299-308.
- [55] A. Baumberg and D.C. Hogg, "Generating spatiotemporal models from training examples," *Image and Vision Computing*, vol. 14, pp. 525-532, 1996.
- [56] A. Pentland and B. Horowitz, "Recovery of non-rigid motion and structure," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 13, pp. 730-742, 1991.

- [57] C. Nastar, "Vibration modes for nonrigid analysis in 3D images," in *European Conference on Computer Vision*, Stockholm, Sweden, 1994, pp. 231–238.
- [58] T.F. Cootes, C.J. Taylor, D. Cooper, and J. Graham, "Training models of shape from sets of examples," in *British Machine Vision Conference*, 1992, pp. 9–18, Springer Verlag.
- [59] A. Hill, A. Thornham, and C.J. Taylor, "Model-based interpretation of 3D medical images," in *British Machine Vision Conference*, Guildford, UK, 1993, pp. 339–349.
- [60] N. Johnson and D.C. Hogg, "Learning the distribution of object trajectories for event recognition," in *British Machine Vision Conference*, Birmingham, UK, 1995, pp. 583–592.
- [61] A. Blake and M. Isard, *Active Contours*, Springer-Verlag, 1998.
- [62] M. Kass, A. Witkin, and D. Terzopoulos, "Snakes: Active contour models," in *IEEE International Conference on Computer Vision*, London, UK, 1987, pp. 259–268.
- [63] A. Baumberg, "Hierarchical shape fitting using an iterated linear filter," in *British Machine Vision Conference*, Edinburgh, Scotland, 1996, pp. 313–323.
- [64] G.D. Sullivan, A. Worrall, and J.M. Ferryman, "Visual object recognition using deformable models of vehicles," in *ICCV Workshop on Context-based Vision*, Cambridge, MA, 1995, pp. 75–86.
- [65] P. Remagnino, A. Baumberg, T. Grove, D. Hogg, A. Worrall, and K. Baker, "An integrated traffic and pedestrian model-based vision system," in *British Machine Vision Conference*, Colchester, UK, 1997, pp. 380–389.
- [66] A. Baumberg and D.C. Hogg, "An adaptive eigenshape model," in *British Machine Vision Conference*, Birmingham, UK, 1995, pp. 87–96.
- [67] C. Wren, A. Azarbayejani, T. Darrell, and A. Pentland, "Pfinder: Real-time tracking of the human body," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 19, pp. 780–785, 1997.
- [68] R.J. Morris and D.C. Hogg, "Statistical models of object interaction," *International Journal of Computer Vision*, vol. 37, pp. 209–215, 2000.
- [69] N. Johnson and D.C. Hogg, "The acquisition and use of interaction behaviour models," in *IEEE Conference on Computer Vision and Pattern Recognition*, Santa Barbara, 1998, pp. 866–871.
- [70] A. Pentland and A. Liu, "Modeling and prediction of human behaviour," *Neural Computation*, vol. 11, pp. 229–242, 1999.
- [71] J. Rittschler and A. Blake, "Classification of human body movement," in *IEEE International Conference on Computer Vision*, 1999, pp. 634–639.
- [72] M. Turk, "Visual interaction with lifelike characters," in *IEEE International Conference on Automatic Face and Gesture Recognition*, Killington, VE, 1996.
- [73] J. MacCormick and A. Blake, "A probabilistic exclusion principle for tracking multiple objects," in *IEEE International Conference on Computer Vision*, Vancouver, Canada, 1999, pp. 572–578.
- [74] J. MacCormick and M. Isard, "Partitioned sampling, articulated objects, and interface-quality hand tracking," in *European Conference on Computer Vision*, Dublin, Ireland, 2000, pp. II:3–19.
- [75] T. Starner and A. Pentland, "Real-time American sign language recognition using Hidden Markov Models," in *International Symposium on Computer Vision*, Coral Gables, FL, 1995, pp. 265–270.
- [76] A. Bobick, S. Intille, J. Davies, F. Baird, C. Pinhanez, L. Campbell, Y. Ivanov, A. Schutte, and A. Wilson, "The Kidsroom: A perceptually-based interactive and immersive story environment," *Presence: Teleoperators and Virtual Environments*, vol. 8, pp. 367–391, 1999.
- [77] M. Scaife and Y. Rogers, "Informing the design of a virtual environment to support learning in children," *International Journal of Human-Computer Studies*, vol. special issue User Centred Design And Implementation Of Virtual Environments, in press.
- [78] A. Pentland, "Smart rooms," *Scientific American*, vol. 274, pp. 54–62, 1996.
- [79] C. Pinhanez and A. Bobick, "Intelligent studios: Modeling space and action to control TV cameras," in *Applications of Artificial Intelligence*, 1997, vol. 11, pp. 285–306.
- [80] D. Beymer, P. McLauchlan, B. Coifman, and J. Malik, "A real-time computer vision system for measuring traffic parameters," in *IEEE Conference on Computer Vision and Pattern Recognition*, San Juan, Puerto Rico, 1997.
- [81] J. Malik, J. Weber, T. Luong, and D. Koller, "Smart cars and smart roads," in *British Machine Vision Conference*, Birmingham, UK, 1995, pp. 367–382.
- [82] VIEWS consortium, "The VIEWS project and wide-area surveillance," in *ESPRIT Workshop at ECCV*, Genoa, Italy, 1992.
- [83] S. Hongeng, F. Bremond, and R. Nevatia, "Bayesian framework for video surveillance applications," in *International Conference on Pattern Recognition*, Barcelona, Spain, 2000.
- [84] S. Hongeng and R. Nevatia, "Multi-agent event recognition," in *IEEE International Conference on Computer Vision*, Vancouver, Canada, 2001.
- [85] R. Howarth, "Interpreting a dynamic and uncertain world: Task-based control," *Artificial Intelligence*, vol. 100, pp. 5–85, 1998.
- [86] R. Howarth and H. Buxton, "Attentional control for visual surveillance," in *ICCV Workshop on Visual Surveillance*, S. Maybank and T. Tan, Eds. 1997, IEEE Press.
- [87] D. Terzopoulos, "Visual modeling for computer animation: Graphics with a vision," *Computer Graphics*, vol. 33, pp. 42–45, 1999.
- [88] D. Terzopoulos and T. Rabie, "Animat vision: Active vision in artificial animals," *Videre: Journal of Computer Vision Research*, vol. 1, pp. 2–19, 1997.
- [89] R. Thibadeau, "Artificial perception of actions," *Cognitive Science*, vol. 10, pp. 117–149, 1986.
- [90] T. Troscianko, A. Holmes, J. Stillman, M. Mirmehdi, and D.B. Wright, "Will they have a fight? The predictability of natural behaviour viewed through CCTV cameras," *Perception*, vol. 30, pp. supplement–72, 2001.
- [91] J.O. Eklundh, P. Nordlund, and T. Uhlin, "Issues in active vision: Attention and cue integration/selection," in *British Machine Vision Conference*, Cambridge, UK, 1996, pp. 1–12.
- [92] M.R. Land and P. McLeod, "From eye movements to actions: How batsmen hit the ball," *Nature Neuroscience*, vol. 3, pp. 1340–1345, 2000.