

Combining generative and discriminative models in a framework for articulated pose estimation

Rómer Rosales*

Probabilistic and Statistical Inference Group
Dept. of Electrical and Computer Engineering
University of Toronto
Toronto, ON M5S 3G4 CANADA
romer@psi.toronto.edu

Stan Sclaroff

Image and Video Computing Group
Dept. of Computer Science
Boston University
Boston, MA 02215 USA
sclaroff@cs.bu.edu

Abstract

A framework is presented that enables reliable estimation of the articulated pose of the human body or human hand, given a single image. Pose estimation is formulated as a statistical inference problem, where the goal is to find a posterior probability distribution over poses. The framework combines two modeling approaches, one discriminative and the other generative. The discriminative model consists of a set of several mapping functions that are estimated automatically from a labeled training set of body poses and their respective image features. A key advantage of the discriminative formulation is that it can model ambiguous, one-to-many mappings that may yield multiple valid articulated pose hypotheses. The generative model is defined in terms of a computer graphics rendering of poses. While this generative model offers an accurate way to relate observed (image features) and hidden (body pose) random variables, it is difficult to use it directly in pose estimation, since inference is intractable. In contrast, inference with the discriminative model is tractable, but considerably less accurate for the problem of interest. A combined discriminative/generative formulation is derived that leverages the complimentary strengths of both models in a principled framework for articulated pose inference. Two efficient pose inference algorithms are derived from this formulation; the first is deterministic and the second non-deterministic. Performance of the framework is quantitatively evaluated in estimating articulated pose of both the human hand and human body.

Keywords: Human body pose, hand pose, nonrigid and articulated pose estimation, statistical inference, generative and discriminative models, mixture models, Expectation Maximization algorithm.

*Rómer Rosales is now at MIT Computer Science and Artificial Intelligence Laboratory, email romer@csail.mit.edu

1 Introduction

An essential task for vision systems is to infer the state of the world given some form of visual observations. From a computational perspective, this typically involves facing an ill-posed problem; relevant information is lost via projection of the three-dimensional world into a two-dimensional image. In this paper, the focus is on inferring the pose of an articulated object in an image, in particular the pose of a human body or human hand. Humans can often solve such pose inference problems, even when given only a relatively poor-quality, low-resolution, monocular image. It is believed that humans employ extensive prior knowledge about human body structure and motion in solving this ill-posed task [21]. In this paper, we consider how a computer vision system might learn such knowledge in the form of probabilistic models, and how to employ such models in an algorithm for reliable pose inference.

For purposes of computation, the inference task can be defined as follows: given an observation vector $\mathbf{x} \in \mathcal{R}^c$ that was extracted from an image of a person, infer the parameterized articulated pose as a vector $\mathbf{h} \in \mathcal{R}^t$. The vector spaces \mathcal{R}^c and \mathcal{R}^t are continuous. In a generic machine learning framework, pose inference might be regarded as a function $\varphi : \mathcal{R}^c \rightarrow \mathcal{R}^t$ that maps an input vector of visual observations to an output vector describing the most likely articulated pose. More generally, the mapping function could produce a posterior probability distribution, $\varphi : \mathcal{R}^c \rightarrow \mathcal{P}$, where \mathcal{P} is a family of probability density functions on \mathcal{R}^t . A number of general questions arise. What type or form should the mapping function φ take? How can the mapping function be estimated from training data? How can the approach incorporate prior knowledge about the problem structure? How can inference be performed efficiently or approximately if exact inference is intractable? These questions are fundamental and common in statistical learning, and only in certain basic problem instances [30] are the answers immediately clear.

Several perspectives or viewpoints could be taken to approach learning tasks. However, often it is not clear which seems more suitable for the problem at hand. It will be useful for the purpose of this paper to distinguish two major viewpoints: the generative and the discriminative viewpoints (*e.g.*, see [27]). In the case of learning generative models, a joint distribution $p(\mathbf{x}, \mathbf{h})$ over the random variables of the model (here simply \mathbf{h} and \mathbf{x}) is estimated from data. Then, given an observation, *e.g.*, \mathbf{x} , a posterior probability $p(\mathbf{h}|\mathbf{x})$ over the unobserved random variables could, in theory¹, be calculated by invoking Bayes' rule. In contrast, using discriminative models the posterior distribution $q(\mathbf{h}|\mathbf{x})$ is directly learned or estimated (see *e.g.*, [25, 36, 27] for further comparisons between these viewpoints). In this paper, we favor the idea that for pose estimation, the advantages of each of these viewpoints could be exploited in a single framework.

¹However, in practice this task can be intractable or lack analytic solutions. This is an important open problem in statistics.

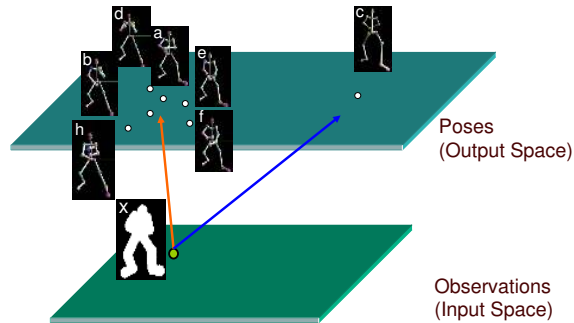


Figure 1: Example ambiguity in mapping body silhouette cues in \mathcal{R}^c to articulated body poses in \mathcal{R}^t . Given silhouette \mathbf{x} , poses \mathbf{a} – \mathbf{h} are all valid hypotheses. In general, entire regions in \mathcal{R}^t may contain valid poses.

If we try to learn a mapping directly, let us say by estimating the parameters of a parameterized function $\phi : \mathcal{R}^c \rightarrow \mathcal{R}^t$ as in a discriminative approach, we encounter several problems. The form required for ϕ may not be simple, because the mapping from observations to articulated poses is generally ambiguous (one-to-many), and therefore no single function can perform this mapping. An example is illustrated in Fig. 1. The arm locations cannot be uniquely inferred given the silhouette \mathbf{x} and therefore, \mathbf{a} – \mathbf{h} are all plausible pose configurations. The hands and arms can move in such a way that the silhouette does not change. Note also that pose \mathbf{c} is the reflection of \mathbf{a} : the camera looks at the back rather than at the front of the body. There may be different regions in \mathcal{R}^t that correspond to ranges of valid poses, and these regions may not be connected, e.g., some viewed from the front and others from behind.

Let us now consider the inverse problem: given an articulated pose vector \mathbf{a} , generate its silhouette \mathbf{x} . With a good computer graphics model of the human body, one can easily render the silhouette \mathbf{x} . Thus, we can define a very accurate generative model of what we refer to as the inverse mapping of pose parameters to image features $\zeta : \mathcal{R}^t \rightarrow \mathcal{R}^c$. This is a key part of our problem definition and it will play an important role in developing the framework presented in this paper. The inverse mapping ζ provides very useful information about the structure of the problem; however, it cannot be incorporated straightforwardly in a discriminative approach. Despite the simplicity of ζ , its inverse may still be complex or not even exist.

In summary, the one-to-many nature of the problem of mapping image features to body poses precludes the use of discriminative supervised learning methods that fit a single (or finite number of) functions to the data, e.g., most neural networks, support vector machines, simple least-squares, boosting, etc. On the other hand, we have access to the *inverse* $\zeta : \mathcal{R}^t \rightarrow \mathcal{R}^c$ that given a body pose can produce the corresponding image features, which can be used to define a very accurate generative model. However, as will be shown




Figure 2: Simplified schematic illustration behind our method for the case of inferring body pose: (a) Given an input vector \mathbf{x} , we generate a set of hypotheses. (b) The inverse mapping function ζ is employed in evaluating each hypothesis.

later, this accurate generative model is challenging to use directly in body pose inference. The view taken in this paper is that it can be effective to use the individual advantages of these two complimentary approaches (discriminative and generative) to formulate an efficient solution to the pose inference/learning problem.

2 Related Work

In computer vision, recovery of articulated body pose from images is often formulated as a *tracking* problem. Usually, link-joint models comprised of 2D or 3D geometric primitives are designed beforehand to roughly match the specific morphology of the target in question [6, 10, 13, 28, 31, 37, 11, 39]. Mesh models have also been used as an alternative to link-joint models [15]. At each frame, these geometric models are fitted to the image to minimize some cost function that favors the overlap of the model and associated image regions (or motion). Although usually not stated, the fitting or cost function in many cases implicitly defines (or can be used to define) a generative model of the observed image. Despite their descriptive power, this family of approaches has a number of critical drawbacks. Generally, a non-linear optimization problem must be solved at every frame (sometimes equivalent to inference in a complex generative model). Careful manual placement of the model on the first frame in a video sequence is also required. Moreover, tracking in subsequent frames tends to be sensitive to errors in initialization and numerical drift; as a result, these systems cannot recover from tracking errors in the middle of a sequence.

To address these weaknesses, specialized dynamical models have been proposed [20, 28, 29]. These methods learn a prior distribution over some specific motion class, such as walking. This prior is used to predict and hopefully improve the pose estimates in future frames. However, this strong prior substantially limits the generality of the motions that can be tracked; a prior for a given class of motions is generally useless when used for tracking objects undergoing a different class of motion, e.g., walking vs. dancing.

Other methods for constrained tracking include [4], where a subspace of allowable motions is learned from a set of examples. These examples and the model (usually linear) are hoped to be sufficient to span the set of possible motions to be seen during tracking. Thus, pose inference involves finding a linear projection of the observed data onto the motion subspace. This subspace approach enforces a strong prior; as mentioned previously, this limits the generalization of the model to classes of motions not seen in the training set. Furthermore, articulated motion is generally non-linear, and cannot be easily explained as a linear projection.

In our approach we avoid matching image features (e.g., image regions, points, or articulated models) from frame to frame. Therefore, we do not refer to our approach as *tracking*, per se. This is in direct contrast with the techniques mentioned above. A number of other approaches also depart from the aforementioned tracking paradigm. We summarize these next.

In [18] a statistical approach is employed in reconstructing the 3D motions of a human figure. The approach employs a Gaussian probability model for short human motion sequences. It is assumed that 2D tracking of the joint positions in the image is given; therefore, this assumption implicitly incurs the restrictions found in all tracking approaches.

In [40] dynamic programming is used to calculate the best global matching of image points to predefined body joints, given a learned probability density function of the position and velocity of body features. Although not explicitly mentioned by the authors, the probability function is defined by a triangulated acyclic graph. Thus, inference is feasible due to the running intersection property [22, 30]. Still, in this approach, the image points and model initialization must be provided by hand or through some other method.

In [5], the manifold of human body dynamics is modeled via a hidden Markov model with an entropic prior. Once the states are inferred from observations, a quadratic cost function is used to generate a continuous path in configuration space, *i.e.*, body pose space.

In all of the non-tracking approaches just referred, models of *motion* were estimated from data. Although the approach presented in this paper can be used to model dynamics, we argue that when general human motion dynamics are to be learned, the amount of training data, model complexity, and computational resources required are impractical. As a consequence, models with unacceptably large priors towards specific motions are generated. Although by not modeling the dynamics we may be ignoring information that could be used to further constrain the inference process, there are some benefits. For instance, a model for inferring body pose that does not consider dynamics provides invariance with respect to speed (*i.e.*, sampling differences) and direction in which motions are performed. This happens simply because this model treats configurations as temporally independent of each other. Other approaches that use a single image include [3, 14, 41]; however, most of these methods also require that projected joint locations be given as

number of training examples	N
training set	$\mathcal{Z} = \{\mathbf{z}_1, \dots, \mathbf{z}_N\}$
training example (input,output) pair	$\mathbf{z}_i = (v_i, \psi_i)$
input (feature) training vector	$v_i \in \mathfrak{R}^c$
output (pose) training vector	$\psi_i \in \mathfrak{R}^t$
generative and discriminative models probability distributions	p, q (respectively)
observation random variable (<i>e.g.</i> , image moments)	$\mathbf{x} \in \mathfrak{R}^c$
hidden random variable of pose parameters	$\mathbf{h} \in \mathfrak{R}^t$
feedback (rendering) function (for generative model)	$\zeta : \mathfrak{R}^t \rightarrow \mathfrak{R}^c$
number of samples during inference	S
a particular observation or input image feature	\mathbf{x}^*
output (pose) hypothesis (a sample from $q(\mathbf{h} \mathbf{x}^*)$)	\mathbf{h}_k
estimate of most likely output hypothesis	$\hat{\mathbf{h}}$
discrete set of labels for mixture components	$\mathcal{C} = \{1, \dots, M\}$
hidden random variables assigning mixture component to training samples	$\mathbf{y} = (y_1, \dots, y_N), y_i \in \mathcal{C}$
prior probability that mixture component k will be used	$\lambda_k = Q(y = k)$
mapping function parameter vector	θ_k
discriminative model parameters (to be learned)	$\theta = (\theta_1, \dots, \theta_M, \lambda_1, \dots, \lambda_M)$
posterior probability of k -th mixture component for \mathbf{z}_i during EM	$\tilde{Q}(y_i = k) = Q(y_i = k \psi_i, v_i, \theta)$

Table 1: Some mathematical symbols used in this paper.

input. In our approach this is not necessary.

Our approach can be thought of mapping visual features to likely body configurations. Following a machine learning paradigm, stochastic functions that map visual features to pose parameters are approximated from training data. A unique aspect of our approach is the combined use of (1) these mapping functions (defining a discriminative model) with (2) the inverse mapping function ζ (defining a generative model). Generally speaking, after multiple poses have been inferred from just the visual cues, ζ transforms these pose configurations back to the visual cue (observation) space. In this space, we can then automatically *choose* among a set of reconstruction hypotheses. This is a fully probabilistic inference process. Our approach avoids the need for manual initialization or tracking; it thereby avoids the consequent disadvantages of tracking. Remarkably, relatively few computations are required for inference. We will now formalize and explain our approach in detail.

3 Probabilistic Models

We propose to a probabilistic, nonlinear framework for combining generative and discriminative models for articulated pose estimation. The framework employs a set of M functions $\phi_k : \mathfrak{R}^c \rightarrow \mathfrak{R}^t$, each associated to a mixture component in a mixture distribution. Each function maps certain sub-domains of the input

space (cues) onto the output space (poses). These functions are estimated automatically from training data via a variant of the Expectation-Maximization algorithm. The learned conditional distribution is then used as an approximation to an accurate (generative model) distribution defined using the inverse function ζ , for which inference is intractable. This basic idea is shown in a schematic way in Fig. 1. The approximation is employed in a similar way as a proposal distribution is used to approximate sampling from a more complex distribution.

We begin by formally defining both the discriminative and generative models to be employed. The discriminative model will be estimated from training data and the generative model will be defined by a rendering function ζ . These models represent two views of the same problem and will be used together in our framework.

3.1 The Discriminative Model

Let $\mathcal{Z} = \{\mathbf{z}_1, \dots, \mathbf{z}_N\}$ be an observed training set of input-output pairs $\mathbf{z}_i = (v_i, \psi_i)$. Each $v_i \in \mathbb{R}^c$ is an input (feature) vector, and each $\psi_i \in \mathbb{R}^t$ is its corresponding output (pose) vector. A summary of mathematical symbols used in this formulation is provided in Table 1.

We will approach our discriminative problem as one of hidden variable density estimation. We begin by introducing the unobserved random variable $\mathbf{y} = (y_1, \dots, y_N)$. In our model any y_i has as its domain the discrete set $\mathcal{C} = \{1, \dots, M\}$ of labels for the specialized mapping functions, and can be thought of as the function number used to map the i -th training pair, \mathbf{z}_i . Thus M is the number of specialized mapping functions. Our model uses parameters $\theta = (\theta_1, \dots, \theta_M, \lambda_1, \dots, \lambda_M)$, where θ_k represents the parameters of the k -th mapping function, and is the prior probability that the k -th mapping function will be used to map an input-output pair.

Using a maximum-likelihood criterion, we are interested in finding the optimal parameter settings for our model; thus, we seek to maximize the joint log-probability:

$$\theta^* = \arg \max_{\theta} \log q(\mathcal{Z}|\theta). \quad (1)$$

Assuming independence of observations given θ , we obtain:

$$\theta^* = \arg \max_{\theta} \sum_i \log q(\mathbf{z}_i|\theta) \quad (2)$$

$$= \arg \max_{\theta} \sum_i \log \sum_k q(\psi_i|v_i, y_i = k, \theta) Q(y_i = k|\theta) q(v_i), \quad (3)$$

where we used the independence assumption $q(v|\theta) = q(v)$. The term $q(v_i)$ describes how input patterns occur. For solving Eq. 3, this term is approximated by the empirical distribution implied by our training data;

as a consequence, patterns that occur more often will have a larger effect in the maximization of Eq. 3. Due to the sum of terms inside the logarithm of Eq. 3, this optimization is generally intractable. However, a variety of practical approximate optimization methods exist, for example, methods that are based on alternating minimizations [8]. Expectation Maximization (EM) [9, 26] updates are described in Sec. 4.

3.1.1 Choice of a Likelihood Function

Note that the above formulation is general. In particular, the form of the probability $q(\psi|v, y, \theta)$ was not specified. A key question in instantiating our approach is: what form should be used for $q(\psi|v, y, \theta)$? This is the probability that output ψ was generated by the mapping function y , given the input v and model parameters θ . In this work we analyze the following possible cases:

1. A Gaussian joint distribution of input-output vectors: $q(v, \psi|y, \theta) = \mathcal{N}((v, \psi); \mu_y, \Sigma_y)$.
2. A Gaussian distribution, whose mean is the output of the y -th mapping function: $q(\psi|v, y, \theta) = \mathcal{N}(\psi; \phi_y(v, \theta), \Sigma_y)$.

3.2 The Generative Model

Our approach also involves the use of a generative model of images (or image features). In the problem of human body pose estimation from a single image this generative model can be defined in a simple way. We will assume that an image or image features are generated by sampling a pose from a prior distribution $p(\mathbf{h})$ and an image is then generated using the rendering function ζ such that:

$$p(\mathbf{x}|\mathbf{h}) = \mathcal{N}(\mathbf{x}; \zeta(\mathbf{h}), \Sigma_\zeta). \tag{4}$$

It is important to notice that despite the fact that the generative model can be defined in a simple manner, the function ζ is of a complex form. In our case, this makes probabilistic inference intractable as will be further explained later.

4 Learning

An approximation method will be used in learning the discriminative model parameters. We will employ an Expectation Maximization (EM) approach. EM provides a general framework for solving the maximum likelihood parameter estimation problem in statistical models with hidden variables, like Eq. 3. Since the EM algorithm is well known [9, 2, 26], we will only provide derivations specific to our formulation.

Note that the unobserved random variables y_i are independent given \mathbf{z}_i . Thus, the E-step reduces to computing the posterior probabilities for each y_i given the model parameters and observed data. We will denote this posterior $Q(y_i = k|\psi_i, v_i, \theta)$ using the shortcut notation $\tilde{Q}^{(t)}(y_i = k)$. We then have:

$$\tilde{Q}^{(t)}(y_i = k) = \lambda_k q(\psi_i|v_i, y_i = k, \theta^{(t-1)}) / \sum_{j \in \mathcal{C}} \lambda_j q(\psi_i|v_i, y_i = j, \theta^{(t-1)}). \quad (5)$$

Stated differently, this step estimates the responsibility of each mapping function, ϕ_k for each data point, \mathbf{z}_i . $\tilde{Q}^{(t)}(y_i = k)$ represents the so called responsibility of function k for data pair i . Also recall that $\lambda_k = Q(y_i = k)$ is the prior probability that function k be used.

The M-step consists of finding $\theta^{(t)} = \arg \max_{\theta} E_{\tilde{Q}^{(t)}}[\log q(\mathcal{Z}, \mathbf{y}|\theta)]$. In both of our cases we can show that this is equivalent to finding:

$$\theta^{(t)} = \arg \max_{\theta} \sum_i \sum_{k \in \mathcal{C}} \tilde{Q}^{(t)}(y_i = k) [\log q(\mathbf{z}_i|y_i = k, \theta) + \log Q(y_i = k|\theta)]. \quad (6)$$

It is important to mention that this is valid if $q(\mathbf{z}_i|\theta)$ depends on y_i and not on y_j , for any $j \neq i$. Note that for the distributions discussed above, this is true. We now present solutions for the cases described above.

4.1 Case (1)

In this case we have:

$$q(v, \psi|y, \theta) = \mathcal{N}(v, \psi; \mu_y, \Sigma_y) = \mathcal{N}\left(\begin{bmatrix} v \\ \psi \end{bmatrix}; \begin{bmatrix} \mu_v \\ \mu_{\psi} \end{bmatrix}, \begin{bmatrix} \Sigma_{vv} & \Sigma_{v\psi} \\ \Sigma_{v\psi}^{\top} & \Sigma_{\psi\psi} \end{bmatrix}\right)_y, \quad (7)$$

where the subscript y is simply the mapping function number. We can show that the parameter learning problem is reduced to a mixture of Gaussian estimation, for which it is straightforward to estimate θ using EM. Moreover, the Bayesian estimate of ψ given an observed v is also Gaussian: $q(\psi|v, y, \theta) = \mathcal{N}(\psi; \mu_{\psi} + \Sigma_{v\psi}^{\top} \Sigma_{vv}^{-1}(v - \mu_v), \Sigma_{\psi\psi} - \Sigma_{v\psi}^{\top} \Sigma_{vv}^{-1} \Sigma_{v\psi})_y$. Therefore in case (1), each function ϕ_k is just the mean of the conditional distribution

$$\phi_k(v, \theta) = (\mu_{\psi} + \Sigma_{v\psi}^{\top} \Sigma_{vv}^{-1}(v - \mu_v))_{y=k}. \quad (8)$$

The confidence of the estimate is given by the covariance $\Sigma_k = (\Sigma_{\psi\psi} - \Sigma_{v\psi}^{\top} \Sigma_{vv}^{-1} \Sigma_{v\psi})_{y=k}$. However, this expression does not depend on the input, a sometimes undesirable consequence of the given model. Thus, each function ϕ_k is linear in the input vector from \mathfrak{R}^c .

4.2 Case (2)

In this case we have:

$$\frac{\partial E}{\partial \lambda_k} = \sum_i \tilde{Q}^{(t)}(y_i = k) \frac{\partial}{\partial \lambda_k} \log Q(y_i = k | \theta) \quad (9)$$

$$\frac{\partial E}{\partial \Sigma_k} = \sum_i \tilde{Q}^{(t)}(y_i = k) \frac{\partial}{\partial \Sigma_k} \log q(\psi_i | y_i = k, v_i, \theta_k) \quad (10)$$

$$\frac{\partial E}{\partial \theta_k} = \sum_i \tilde{Q}^{(t)}(y_i = k) \left[\left(\frac{\partial}{\partial \theta_k} \phi_k(v_i, \theta_k) \right)^\top \Sigma_k^{-1} (\psi_i - \phi_k(v_i, \theta_k)) \right], \quad (11)$$

where E is the cost function that we would like to maximize in Eq. 6.

This gives the following update rules for λ_k and Σ_k , where Lagrange multipliers were used to incorporate the constraint that the sum of the λ_k 's is 1:

$$\lambda_k^{(t)} = \frac{1}{N} \sum_i \tilde{Q}^{(t)}(y_i = k) \quad (12)$$

$$\Sigma_k^{(t)} = \sum_i \tilde{Q}^{(t)}(y_i = k) (\psi_i - \phi_k(v_i, \theta_k)) (\psi_i - \phi_k(v_i, \theta_k))^\top / \sum_i \tilde{Q}^{(t)}(y_i = k) \quad (13)$$

To keep the formulation general, we have not yet defined the form of the mapping functions ϕ_k . Whether or not we can find a closed form solution for the update of θ_k depends on the form of ϕ_k . For example if ϕ_k is a non-linear function, we may have to use iterative optimization to find $\theta_k^{(t)}$. If ϕ_k yields a quadratic form, then a closed form update exists.

Regarding our generative model, there is very little learning involved. If ζ is very accurate, then we could also tell very accurately the image that will be generated given a body pose \mathbf{h} . In practice ζ can be defined only approximately. We account for this by properly setting Σ_ζ depending of how much noise is expected to be present in the observations. This can also account for inaccuracies in the geometric model.

5 Inference

In this section, we refer to probabilistic inference as finding a full probability distribution for \mathbf{h} given that $\mathbf{x} = \mathbf{x}^*$ once an observation \mathbf{x}^* has been made (*e.g.*, some image features were observed).

5.1 Inference using the Discriminative Model Alone

A valid approach to inference is to use the discriminative model alone. In order to understand how this differs from our proposed solution (where we combine both, generative and discriminative models), we will now

show what inference involves in terms of maximum a posteriori (MAP) estimation using the discriminative model.

In a general sense inference involves finding a full probability distribution for \mathbf{h} given \mathbf{x}^* ; the discriminative model directly provides this expression. In MAP estimation we just have to maximize it (*i.e.*, we want to find the most likely output hypothesis $\mathbf{h} \in \mathfrak{H}$ for a given observation $\mathbf{x}^* \in \mathfrak{R}^c$): $\hat{\mathbf{h}} = \arg \max_{\mathbf{h}} q(\mathbf{h}|\mathbf{x}^*) = \arg \max_{\mathbf{h}} \sum_y q(\mathbf{h}|\mathbf{x}^*, y)Q(y)$, where $q(\mathbf{h}|\mathbf{x}^*)$ is a shorthand for $q(\mathbf{h}|\mathbf{x} = \mathbf{x}^*)$. Any further treatment depends on the properties of the probability distributions involved.

In both Cases (1) and (2) considered in previous sections, we can write $q(\mathbf{h}|\mathbf{x}, y) = \mathcal{N}(\mathbf{h}; \phi_y(\mathbf{x}), \Sigma_y)$. Thus, in either case we have that $q(\mathbf{h}|\mathbf{x}^*)$ is a mixture of Gaussians and if we want to find the MAP estimate we need to solve: $\hat{\mathbf{h}} = \arg \max_{\mathbf{h}} \sum_y \mathcal{N}(\mathbf{h}; \phi_y(\mathbf{x}^*), \Sigma_y)Q(y)$.

This result was obtained from performing (MAP) inference using our learned discriminative model alone, where we learned $q(\mathbf{h}|\mathbf{x})$ as an approximation to the true distribution defined by $p(\mathbf{h}|\mathbf{x})$, using the training data. Even though we could simply adopt this as a solution, it should not be surprising that we could improve upon this by using our knowledge of p , the generative model.

5.2 Inference Using the Generative Model Alone

Using the generative model, inference involves finding the posterior $p(\mathbf{h}|\mathbf{x} = \mathbf{x}^*)$ ($p(\mathbf{h}|\mathbf{x}^*)$ as a shorthand):

$$p(\mathbf{h}|\mathbf{x}^*) = \frac{p(\mathbf{x}^*|\mathbf{h})}{p(\mathbf{x}^*)}p(\mathbf{h}) = \frac{1}{Z_p}\mathcal{N}(\mathbf{x}^*; \zeta(\mathbf{h}), \Sigma_\zeta)p(\mathbf{h}) \quad (14)$$

$$Z_p = \int \mathcal{N}(\mathbf{x}^*; \zeta(\mathbf{h}), \Sigma_\zeta)p(\mathbf{h})d\mathbf{h}. \quad (15)$$

There are however at least two difficult obstacles for achieving this:(1) The integral in Eq. 15 cannot be solved easily and moreover, (2) we do not have an expression for $p(\mathbf{h})$.

In MAP estimation, the goal is to find $\hat{\mathbf{h}}$ such that: $\hat{\mathbf{h}} = \arg \max_{\mathbf{h}} p(\mathbf{h}|\mathbf{x}) = \arg \max_{\mathbf{h}} \mathcal{N}(\mathbf{x}; \zeta(\mathbf{h}), \Sigma_\zeta)p(\mathbf{h})$, needless to say that in the case of body pose estimation, this is in general, a highly complex non-linear optimization problem (tracking) as we have seen before.

A key idea in this paper is that both obstacles would become much simpler if, somehow, we could accurately obtain samples from $p(\mathbf{h}|\mathbf{x})$. Those samples could be used to (1) approximate this posterior and (2) find the most likely sample for MAP estimation. However sampling accurately from a given distribution, in particular $p(\mathbf{h}|\mathbf{x})$, is in general an open problem [24].

5.3 Inference and Importance Sampling. Combining Generative and Discriminative Models

In general, sampling can be used to estimate expectations of a given function $I(w)$ with respect to some probability density $\pi(w)$ that we can evaluate at any point, but that we cannot sample from. Let us say we need to calculate the integral $\mathcal{I} = \int \pi(w)I(w)dw$, by approximating \mathcal{I} employing R samples: $\hat{\mathcal{I}} = \frac{1}{R} \sum_{r=1}^R I(w^{(r)})$.

The question is how to appropriately generate the samples to obtain the best estimate for \mathcal{I} . Sampling can be used in more general tasks. Usually it is only necessary to evaluate $\pi^*(w)$ which is equal to $\pi(w)$ to within a multiplicative factor, $\pi(w) = \pi^*(w)/Z$.

Since we cannot usually generate samples accurately, we need to account for our sampling inaccuracies. Importance sampling is a method that accounts for this as follows: first we come up with a proposal distribution $\pi'(w)$, which we can also evaluate (within a multiplicative factor) but from which it is possible to sample; then we sample from $\pi'(w)$, but also correct for the bias introduced when sampling, obtaining:

$$\hat{\mathcal{I}} = \frac{1}{R} \sum_{r=1}^R \frac{\pi^*(w^{(r)})}{\pi'(w^{(r)})} I(w^{(r)}). \quad (16)$$

It can be shown that when $R \rightarrow \infty$, $\sqrt{R}(\hat{\mathcal{I}} - \mathcal{I}) \sim \mathcal{N}(0, \sigma_{\pi'}^2)$, with: $\sigma_{\pi'}^2 = \int (\frac{\pi^*(w)}{\pi'(w)} I(w) - \mathcal{I})^2 \pi'(w) dw$. Thus, the expected variance of our estimate is proportional to $\sigma_{\pi'}^2$, and inversely proportional to R [24].

We would like to know what the optimal proposal distribution is in order to maximize the accuracy of our estimator (minimize $\sigma_{\pi'}^2$). The optimal proposal distribution to approximate \mathcal{I} is given by a result in [35, 7]:

$$\pi'(w) = \pi^*(w) / \int \pi^*(w) dw, \quad (17)$$

the normalized function $\pi^*(w)$. Since, on purpose, we introduced $\pi^*(w)$ first as a unnormalized distribution originating from $\pi(w)$, we know that upon normalization we will get back the original distribution $\pi(w)$. A similar justification was used in [7], but their way of building a proposal distribution is different from ours.

Thus, in finding a posterior distribution $p(\mathbf{h}|\mathbf{x})$ for body poses given observed image features the partition function in Eq. 15 can be computed as follows:

$$Z_p = \int p(\mathbf{x}^*, \mathbf{h}) d\mathbf{h} \approx \frac{1}{S} \sum_{s=1}^S p(\mathbf{x}^*, \mathbf{h}^{(s)}) / p(\mathbf{h}^{(s)}|\mathbf{x}^*), \quad (18)$$

using importance sampling to obtain the samples $\mathbf{h}^{(s)}$, taken from $p(\mathbf{h}|\mathbf{x}^*)$. We have shown that the best proposal distribution is $p(\mathbf{h}|\mathbf{x}^*)$; unfortunately this is the distribution we are after, and of course we cannot just sample from it.

The main reason behind using generative and discriminative models together is to tackle this particular problem of sampling from a good distribution. We will use the learned distribution $q(\mathbf{h}|\mathbf{x})$ (the discriminative model) to approximate $p(\mathbf{h}|\mathbf{x})$, but just at $\mathbf{x} = \mathbf{x}^*$. This approximation is in terms of maximum likelihood estimation and can also be seen as minimizing the KL divergence between the empirical distribution p_e , given by the training data, and the model distributions q , parameterized by θ :

$$\text{KL}(p_e(\mathbf{x}, \mathbf{h})||q(\mathbf{x}, \mathbf{h})) = \int p_e(\mathbf{x}, \mathbf{h}) \log[p_e(\mathbf{x}, \mathbf{h})/q(\mathbf{x}, \mathbf{h})]d\mathbf{h}d\mathbf{x}, \quad (19)$$

which can be proven to be equivalent to:

$$\arg \min_{\theta} E_{p_e(\mathbf{x})}[\text{KL}(p_e(\mathbf{h}|\mathbf{x})||q(\mathbf{h}|\mathbf{x}))], \quad (20)$$

where θ parameterizes q . In practice, the expectation becomes a sum over the training data pairs, and we obtain Eq. 1. Thus, the optimal distribution in this sense is the one that results from solving Eq. 1, to obtain $q(\mathbf{h}|\mathbf{x})$. Of course, we assume that the data is composed by representative examples from p , so that the empirical distribution p_e is at all useful. Eq. 17 justifies this choice since it tells us that in order to find a good approximation for the posterior $p(\mathbf{h}|\mathbf{x})$ we should find a proposal distribution that is similar to it, as intuitively expected. We may then ask if we could use this proposal distribution alone. The reason why this is not a good idea is that, since we cannot usually find a proposal distribution that matches the true posterior perfectly, using this proposal distribution alone is expected to perform worse than when combined with our accurate generative model. This is mainly because in regions where the proposal distribution q is bad at approximating p , we can always evaluate p and note the error or discrepancy.

The distribution $q(\mathbf{h}|\mathbf{x})$ is an approximation to $p(\mathbf{h}|\mathbf{x})$ in the space of all distributions with the structure specified by the discriminative model (a mixture model in our case). For Gaussian mixture models, it is known that this approximation can be made as accurate as we wish in the limit of infinite data and mixture components. Interestingly, we do not need to know explicitly what $p(\mathbf{h})$ is in our generative model. Of course, this is implicitly specified by the training data. Thus, even if we use a not so good assumption for $p(\mathbf{h})$, still we know what we need to do in order to achieve a good estimate of the posterior. This is helpful since we do not really know accurately what $p(\mathbf{h})$ is (given that we may not have enough data to estimate it accurately). In the following we simply use an uniform distribution (in a reasonable finite domain).

To summarize, in order to compute the posterior distribution of body poses \mathbf{h} , given an observation of image or image features \mathbf{x}^* , we calculate and estimate for $p(\mathbf{h}|\mathbf{x}^*)$ as follows:

$$\hat{p}(\mathbf{h}|\mathbf{x}^*) = \frac{1}{\hat{Z}_p} \mathcal{N}(\mathbf{x}^*; \zeta(\mathbf{h}), \Sigma_{\zeta}) p(\mathbf{h}), \quad (21)$$

with \hat{Z}_p given by Eq. 18, but substituting $p(\mathbf{h}|\mathbf{x}^*)$ by $q(\mathbf{h}|\mathbf{x}^*)$ and using samples from $q(\mathbf{h}|\mathbf{x}^*)$.

5.4 Non-deterministic MAP Estimation: Multiple Samples (MS)

We are usually interested in providing likely samples from the posterior distribution, in particular we might be interested in the most likely \mathbf{h} . This is the idea behind MAP estimation, where we are interested in finding $\hat{\mathbf{h}} = \arg \max_{\mathbf{h}} p(\mathbf{h}|\mathbf{x}^*) = \arg \max_{\mathbf{h}} p(\mathbf{x}|\mathbf{h}^*)p(\mathbf{h})$.

We know that the discriminative model distribution $q(\mathbf{h}|\mathbf{x})$ tries to approximate $p(\mathbf{h}|\mathbf{x})$, and therefore it is good at minimizing the variance of the estimator. Due to this, we will use the discriminative model distribution to provide samples for MAP estimation. In MAP estimation, we sample $\mathcal{H}_{Spl} = \{\mathbf{h}_s\}_{s=1\dots S}$ using the proposal distribution $q(\mathbf{h}|\mathbf{x}^*)$. Given the samples, the problem becomes a discrete optimization problem that can be solved easily:

$$\hat{\mathbf{s}} = \arg \max_s p(\mathbf{x}^*|\mathbf{h}_s) = \arg \min_s (\mathbf{x}^* - \zeta(\mathbf{h}_s))^\top \Sigma_\zeta (\mathbf{x}^* - \zeta(\mathbf{h}_s)), \quad (22)$$

by using the Gaussian form of $p(\mathbf{x}|\mathbf{h})$ as given in Eq. 4. We remark that after using the samples \mathcal{H}_{Spl} as a starting point, other more sophisticated methods could be employed. For example we could use Markov chain Monte Carlo (MCMC) sampling [24, 44] to search for regions of higher probability. Also, instead of stochastic methods, we could employ standard gradient descent methods to locally search for more likely poses \mathbf{h} (as in tracking). These methods may be helpful for some distributions but in general have several drawbacks: (1) They are usually very slow in high dimensions and (2) given finite time, they are not very useful/accurate if the posterior probability is very complex. Some methods have been proposed to alleviate these problems, but this goes beyond our current contribution. Keeping this extension in mind, in this paper we simply use the original samples \mathcal{H}_{Spl} to search for a MAP estimate. These estimates proved to be sufficiently accurate during our experiments.

5.5 Deterministic MAP Estimation: Mean Output (MO)

In certain applications, it might be advantageous to count with a very fast method for computing MAP estimates. Two examples are: when working with multiple articulated bodies and in dynamic settings where it is necessary to provide estimates at a high rate. Even though the time complexity of MS scales linearly with the number of samples, this might not be fast enough. Motivated by speed constraints, here we propose a very fast MAP estimation algorithm that still performs well in experiments. Unlike MS, this algorithm is deterministic.

The structure of the problem, as well as the form of the discriminative distribution components (*i.e.*, conditioned on the mixture label) $q(\mathbf{h}|\mathbf{x}, y)$ employed (Gaussian), make it possible to construct this deterministic approximation. The basic intuition is straightforward. For a given $\mathbf{x} = \mathbf{x}^*$, we ask each mapping

function ϕ_k to give its most likely estimate for \mathbf{h} . We then evaluate the probability of each function’s estimate via the generative model distribution $p(\mathbf{x}^*|\mathbf{h})$. This approximation is good in practice, as will be demonstrated in the experiments.

To justify this deterministic approximation, we note that due to its concavity properties, the probability of the mean is maximal in a Gaussian distribution; *i.e.*, it is the most-likely value. Formally, in both Case (1) and Case (2) described earlier, $q(E[\mathbf{h}|\mathbf{x}^*, y, \theta]) \geq q(\mathbf{h}'|\mathbf{x}^*, y, \theta)$, for any \mathbf{h}' . Consider again the set of samples $\mathcal{H}_{Spl} = \{\mathbf{h}_s\}_{s=1\dots S}$ generated in the MS approximation. We can build a set of samples $\mathcal{H}_\phi = \{\mathbf{h}_k^\phi\}_{k=1\dots M}$ that has the property $\forall y, \max_k q(\mathbf{h}_k^\phi|\mathbf{x}^*, y) \geq \max_s q(\mathbf{h}_s|\mathbf{x}^*, y)$, simply by setting $\mathbf{h}_k^\phi = \phi_k(\mathbf{x}^*, \theta)$.

This insight leads to a deterministic approximation for inference, the *Mean Output* solution (MO). This approximate solution relies on the observation that by considering the means $\phi_s(\mathbf{x}^*)$, we would be considering the most likely output of each mapping function (*i.e.*, each mixture component in the discriminative model), given the input. Obviously we expect the discriminative model provides a good approximation of our generative model posterior distribution as discussed above. Also, the smaller the overlap among the distributions associated with each function, the better the accuracy of this approximation.

In MO approximate inference, the expression to be minimized is the same as that used in Eq. 22, except for the use of the M means instead of the S samples:

$$\hat{k} = \arg \max_{k \in \mathcal{C}} p(\mathbf{x}^*|\mathbf{h}_k^\phi) = \arg \min_{k \in \mathcal{C}} (\mathbf{x}^* - \zeta(\mathbf{h}_k^\phi))^\top \Sigma_\zeta (\mathbf{x}^* - \zeta(\mathbf{h}_k^\phi)). \quad (23)$$

This generally requires substantially less computation than would be required in the MS approach.

6 Example Application: Articulated Pose from Visual Features

The formulation presented in this paper is rather general, and could be applied in a number of supervised learning problems for which the output-to-input (feedback) map is relatively easy to compute; thus allowing us to specify an accurate generative model. To demonstrate and test our framework, we have developed a system that uses the our approach to infer articulated pose from low-level visual features. In particular, we focused on pose estimation of the human hand and body from an image silhouette. In this class of computer vision applications, ground truth datasets for training can be obtained via motion capture gloves or body suits, and computer graphics rendering can be used to generate the input-output pairs used in supervised learning. We will now give details of this demonstration system.




Figure 3: Example of the 86 silhouettes obtained via computer graphics rendering for a given a 3D hand pose. Views are distributed approximately uniformly over the view sphere.

6.1 3D Hand Pose Estimation

The goal is to recover detailed 3D hand pose from silhouette features computed from a single color image. Hand pose is defined in terms of the hand joint angles. In general, we are also interested in global orientation of the hand. We explore two applications: estimation of the internal joint angles only, and later, estimation of both internal joint angles and global orientation of the hand.

6.1.1 Hand Model

We utilize the hand model provided in the VirtualHand programming library [42]. The model parameters are 22 joint angles. For the index, middle, ring and pinky finger, there is an angle for each of the distal, proximal and metacarpophalangeal joints. For the thumb, there is an inner joint angle, an outer joint angle and two angles for the trapeziometacarpal joint. There are also abduction angles between the following pairs of successive fingers: index/middle, middle/ring and ring/pinky. Finally, there is an angle for the palm arch, an angle measuring wrist flexion and an angle measuring the wrist bending towards the pinky finger. However, because the former two wrist angles also encode global orientation, we decided not to model them in our application. Hence, ignoring these two angles, our model has 20 DOF for the internal hand configuration.

All of these 20 angles are relative to two global orientation angles. These two angles will encode the camera viewpoint (or alternatively hand 3D rotation). Imagine a sphere surrounding the hand model, *i.e.*, a fixed hand center point is at the center of the sphere. For ease of reference, we will employ the widely used latitude and longitude notions. The first angle β_1 represents the latitude from which we are looking at the hand, the second angle β_2 represents the longitude. We have defined $\beta_1 \in [0, \pi]$, with zero and π being the *poles* of the sphere and $\beta_2 \in [0, 2\pi)$. Thus, in summary our full hand model has 22 DOF.

6.1.2 3D Hand Motion Datasets

Using a CyberGlove, we collected approximately 9,000 examples of 3D hand poses. This data included hand configurations from American Sign Language (ASL) and other configurations informally performed by several subjects. Using computer graphics and an artificial hand model, we then rendered each captured hand pose from multiple viewpoints on the view sphere. We defined a set of 86 viewpoint angle pairs (β_1, β_2) so that the sphere surface is sampled approximately uniformly. Thus we obtained a full dataset of $9,000 \times 86$ views. Each view has an associated binary image mask (silhouette), and a 22 DOF pose vector. Fig. 3 shows the 86 viewpoints used in the dataset for a particular configuration.

From these silhouettes, we extract the visual features that will be used for further processing. In our implementation, we used two classes of features (these features are not used together): Hu moments and Alt moments. Alt moments [1] are translation and scale invariant, but not rotation invariant. Hu moments [19] are invariant to translation and scaling, but also invariant to rotation in the image plane. These moment features were used in our implementation because they are relatively easy to compute, and they provide invariants that are appropriate for our demonstration application. However, our general formulation can be used with other visual feature representations if desired. Detailed examination of the feature selection problem is outside the scope of this paper, and remains a topic for future research.

We define two experimental datasets:

1. *Hand-Single-View*: In this dataset, the hand is viewed from only one viewpoint $(\beta_1 = \pi/2, \beta_2 = 0)$, generally making the palm of the hand visible. Silhouette features are computed using Alt moments. This yields approximately 9,000 input-output pairs.
2. *Hand-All-Views*: In this dataset, the hand is viewed from all 86 viewpoints. Silhouette features are computed using Hu moments. This yields approximately 750,000 input-output pairs.

6.1.3 Hand Detection and Segmentation

For live video input, we will use video sequences collected with a color digital camera. It will be assumed that these sequences have a static background and only one person is present. In this implementation, we are not considering hand occlusion analysis, which by itself is a difficult task. Our system tracks both hands of the user automatically using a skin color tracker [38, 34].

6.2 2D Human Body Pose Estimation

In this application, our goal is to recover the articulated pose of a human body observed in a single image. The methodology followed is very similar to that used in the estimation of hand pose. However, instead of joint angles, body pose will be specified in terms of marker positions at a predetermined set of joints. We will estimate the 2D positions of these body markers in the image plane.

6.2.1 Human Body Model

The human body model is defined in terms of 20 3D marker positions (60 DOF). The 20 markers are distributed as follows: three markers for the head, three markers for the hip/back bone articulation, plus one marker for each shoulder, elbow, wrist, hand, knee, ankle, and foot. For computer graphics rendering, the body model is composed of cylinders of equal width. The cylinders connect the markers to form the standard human body structure. The thorax is modeled using a wider cylinder. Because we are only interested in the shape of the projected model, we do not include texture or illumination in our rendering.

6.2.2 Human Body Pose Dataset

Human body motion capture data was obtained from several sources: <http://www.biovision.com>, Matt Brand's dataset [5], and several demo sequences in the software package *Character Studio*. In total there are 32 captured sequences that depict variations of different activities: dancing, walking, kicking, waving, throwing, jumping, signaling, crouching down. The total number of frames collected is approximately 7,000, mostly at 30 frames/second. Using computer graphics and our artificial body model, we then rendered each frame from 16 equally-spaced viewpoints on the equator of the view sphere centered at the hip of the body model. For each view, we also used the camera model to obtain the 2D marker positions in the image plane. Thus we obtained a full dataset of approximately $7,000 \times 16$ views. Each view has an associated binary image mask (silhouette) and a 40 DOF projected marker vector. From the silhouettes, we extract the visual features that will be used as input. We have chosen Alt moments [1] as our visual features, mainly due to their ease of computation and invariance to translation and scaling. We call this the *Body-All-Views* dataset.

6.2.3 Detection and Segmentation

For live video input, we use sequences collected with a color digital camera. It is assumed that these sequences have a static background, only one person is present, and the person is fully-visible. We use a simple

and widely-used human body segmentation scheme [17, 43]. The technique employs statistical learning to acquire a model of the background appearance, where each pixel’s color (luminance) is represented by a Gaussian distribution. Segmentation is then approached using maximum-likelihood, where each pixel is classified as belonging to the background or the foreground (human body).

6.3 Common Implementation Details

We now briefly discuss implementation details common to both applications.

6.3.1 Mapping Functions

In Sec. 3, it was not specified what class of (deterministic) mapping functions ϕ_k were to be used. Our framework is practically independent of this choice. However, from Eq. 11 we can notice that there are clear advantages in the M-step if these functions are differentiable with respect to their parameters. In the case of quadratic or linear functions, the M-step can be performed exactly in one step. However, the power of these functions is limited. In our implementation each function is a multi-layer perceptron with one hidden layer (MLP). For this non-linear function there does not exist a closed-form solution for Eq. 11. We used four to five iterations of the conjugate gradient descent method per M-step.

6.3.2 Generative Model Details: Feedback Functions

There are at least two ways to define this function. On the one hand, ζ could be a computer graphics rendering function. On the other hand, we could estimate an approximate $\hat{\zeta}$ given a set of output-input training examples. In our implementation, we experimented with both ideas. For ζ , we used computer graphics renderings of our hand and body models obtained via OpenGL. For $\hat{\zeta}$, we used a one hidden-layer perceptron, with twenty hidden nodes. In our experience, this provides an adequate and efficient approximation.

The approximate feedback function is useful primarily because it is faster to compute than a graphical rendering followed by visual feature computation. The key issue to keep in mind is that the feedback mapping is assumed to be simple (one-to-one or even many-to-one) or that it has a known form, otherwise if we assume too simple functional forms, we would only introduce more estimation errors. Of course, this is just a practical issue. If the feedback mapping is too complex to approximate easily, we could always rely on the available feedback function ζ .

6.3.3 Computational Performance

For an Athlon 1400 PC with 2GB memory, running unoptimized Matlab 6.0 code, it takes approximately five hours to train a model with 10 dimensions (input) and 10 dimensions (output), using 4500 patterns, and 40 single hidden layer perceptrons with five hidden nodes each. The system can infer body poses at approximately 11 frames per second, using the Mean Output (MO) algorithm. This approach's related computations take approximately 70% of this time. This time includes OpenGL-based rendering of body poses in ζ . The rest is spent in segmentation and feature calculations. The Multiple Sample (MS) algorithm takes time proportional to the number of samples used. Of course, segmentation and feature computation for the segmented image is done only once. We noticed that for our implementation, if we use the approximate feedback function, $\hat{\zeta}$, the rendering time is reduced to approximately one-fourth.

6.3.4 Early Stopping During Training

During model training, we used cross-validation for early stopping and to avoid over-fitting as follows:

- *Training data:* Stop if the log-likelihood changes less than 0.5% averaged over the last ten iterations.
- *Held out data:* Stop if the held out data log-likelihood average change is negative over the last ten iterations. Held out data was chosen in the same way as the training and test data.
- *Number of iterations:* Stop if a maximum of 200 iterations is reached.

7 Experimental Results

We now present experimental results obtained using our approach in estimating the pose of the human hand and body. For many additional performance experiments not included due to space limitations, the reader is referred to [33] and for several MO estimation videos to <http://www.psi.toronto.edu/~romer/SMAHandVideos.htm>. The SMA application independent Matlab code can be found at <http://www.psi.toronto.edu/~romer/SMACode.htm>.

7.1 Hand Pose Estimation Given a Fixed Camera Viewpoint

In our first experiments, our approach is tested in the task of recovering 3D human hand pose given a fixed camera viewpoint: a view towards the palm of the hand. For training, we used the *Hand-Single-View* dataset, which contains a total of approximately 9,000 examples. Of these, 3,000 were used for training and the rest for testing. All experiments were performed on a test set that shared no common poses with the training set.

	MO-MAP ($\hat{\zeta}$)	MS-MAP ($\hat{\zeta}$)	MS-20 ($\hat{\zeta}$)	MO-MAP (ζ)	MS-MAP (ζ)	MS-20 (ζ)	Rand/train	Range
$\hat{\mathcal{E}}$	0.1322	0.1667	0.1465	0.1651	0.1769	0.1785	0.4294	1.55
$\sigma_{\hat{\mathcal{E}}}^2$	0.0317	0.0415	0.0371	0.0425	0.0452	0.0547	0.1630	-

Table 2: Mean absolute error $\hat{\mathcal{E}}$ and variance $\sigma_{\hat{\mathcal{E}}}^2$. Inference performance using different rendering functions (ζ and $\hat{\zeta}$) and inference algorithms (MO-MAP and MS-MAP). Also shown, the accuracy of the most probable reconstructions given by MS (MS-20). As a point of comparison, results are presented for an algorithm that randomly chooses one of the training examples as result (Rand/train). The average range of the data is also shown as a reference point. All units are in radians.

The input-output pairs were then defined as follows. The input consisted of 10 Alt moments computed from the silhouette of the hand, as described in Sec. 6.1. The output consisted of 20 joint angles of a human hand linearly encoded by nine values using Principal Component Analysis (PCA).

The number of mixture components for the discriminative model (mapping functions) was set to 20. This number was found to be optimal in the sense of the Minimum Description Length (MDL) principle [32]; we found this number via a rough model search (testing MDL and getting the score for the optimized model with 10,12,...,24 functions). Each mapping function (for each of the Gaussians in the mixture) was a MLP with seven hidden neurons.

7.1.1 Quantitative Results

We randomly selected approximately 4,000 frames not included in the training set. Since ground-truth is available, we used the average absolute difference per joint angle (between ground-truth and estimate) as error measure. Table 2 summarizes our results (see caption).

These experiments quantitatively confirmed that MO inference provides a reasonable approximation, at least for this dataset. Recall from Sec. 5.5 that MO inference was based on the premise that the most-likely reconstruction given by each discriminative mixture component provides a good approximation to the best solution given by the full probability distribution.

Fig. 4 shows example reconstructions obtained via the MO approach. In many cases, the reconstruction is close to the ground truth. In other cases, the silhouette is highly ambiguous, and the reconstruction does not match ground truth. A good example is shown in image pair number 34 (the last row-pair, fourth column), where the camera’s image plane is perpendicular with the axis of the pinky finger. Note that the estimated hand pose disagrees with the ground-truth in the several joint angles associated with this finger. Similar effects with other joint angles can be seen in example pairs 8, 16, 27, etc.

Ambiguous configurations are indeed very common with a binary image representation. Note that in other ambiguous cases shown in Fig. 4 reconstruction is closer to ground truth, *e.g.*, pairs 19, 20, etc.




Figure 4: 40 examples of estimated hand poses chosen uniformly at random. Reconstruction found using the Mean Output (MO) approach. The feedback function used was estimated from data. Each example consists of a pair of images: ground-truth (top), and estimate obtained using the mean output algorithm (bottom).

Possible reasons for this agreement are diverse:

1. The input is not really ambiguous (probabilistically speaking) in the observation space. The other possible outputs (geometrically speaking) associated with this input may be very unlikely given the training set. This depends on the underlying structure of the configuration manifold. One of the main goals of a learning algorithm is to find this structure. Indeed these results show that our algorithm is finding this structure, since in most cases, MO finds a valid sample from the manifold.
2. The learned discriminative model was accurate at modeling the given input using a single mixture component (*i.e.*, few mapping functions were trained to map this input, therefore the rest of the functions produced irrelevant (bad) outputs).
3. By chance, among many very similarly probable solutions, the *right* one was chosen. Of course, even with the help of chance in this case, the discriminative model needed to be accurate enough at approximating the true posterior so that samples were relevant at all.

7.1.2 Experiments with Real Images

We now test our approach using uncalibrated video sequences, where the camera is pointing towards the palm of a person’s hand. On average, the hand occupied an area of approximately 200×200 pixels. Segmentation was obtained as described in Sec. 6.1.3. In the first experiment, we use the MO approach to obtain a single *best* estimate for each segmented hand. Estimates for 40 frames, taken 0.9 seconds apart, are shown in Fig. 5. Visually we can notice that in most cases the estimate is a plausible explanation of the segmented silhouette. However, there are also a few inaccurate reconstructions.

In general, it is expected that the model cannot perform well in all configurations (this is true for almost any machine learning model) due to the following reasons:

1. The proposal distribution $q(\mathbf{h}|\mathbf{x})$ does not resemble the true posterior distribution $p(\mathbf{h}|\mathbf{x})$ at the particular $\mathbf{x} = \mathbf{x}^*$: learning is the result of optimizing an *expected* or average error.
2. The real hand and synthetic hand model features are similar but not the same. Anthropometric differences can influence inference accuracy.
3. Even the best model could fail in some configurations. Information theory tells us that this is always the case except when the *information* in the features is equal to the entropy of the body pose configurations; in other words, when features tell us everything needed about the configuration. Otherwise, there might be multiple explanations for a given visual feature vector.

In order to test the ability of the system to provide these multiple explanations, we tested the Multiple Samples (MS) approach. Fig. 6 shows the estimates found using MS. These estimates can be interpreted as possible hypotheses of hand configurations given the silhouettes. Note that MS tends to bias the hypotheses towards samples from the distribution $q(\mathbf{h}|\mathbf{x}^*)$, but we can account for this when building a full probability distribution, as explained in Sec. 5.3

7.2 3D Hand Pose Reconstruction Given an Unrestricted Camera Viewpoint

Our approach is now tested in the task of recovering 3D human hand pose from an unknown camera viewpoint. For training, we used the *Hand-All-Views* dataset, which contains a total of approximately 750,000 examples. Of these, 18,000 were used for training and the rest for testing. The input-output pairs were then defined as follows. The input consisted of seven Hu moments computed from the silhouette of the hand, as described in Sec. 6.1. The output consisted of 20 internal joint angles of the hand and two orientation angles. This 22 DOF representation was linearly encoded by nine values using PCA.




Figure 5: 40 examples of estimated hand poses captured every 0.9 secs from real video (RV). Reconstruction found using the Mean Output (MO) approach. The feedback function was computed using computer graphics rendering.

The number of mixture components (mapping functions) was set to 45. This number was determined via the MDL criterion, as before (testing for the best MDL score using a model with 35,37,...,51 functions). Each function was a MLP with seven hidden nodes.

7.2.1 Quantitative Results

As before, we computed the absolute error in estimating hand pose, and quantitatively compared this measure across views. Fig. 7 shows the error of the most likely estimate found using the MO approach. From the graphs we see that views towards the palm of the hand (90°) are slightly easier to reconstruct on average, while the variance seems similar across views. As expected, the average error is higher than that obtained for the fixed view hand pose reconstruction experiments. It seems that for unrestricted hand views it is a bit advantageous to use the computer graphics feedback function ζ . This is probably because estimating this inverse mapping $\hat{\zeta}$ (to define the generative model) over unrestricted viewpoint is more complicated than for only frontal hand views (and the mapping is likely to be more complex also).

Fig. 8 shows the results using the MS approach. Fig. 8(a) shows the error associated with the best sample. This error behaves very similarly to the MO error. Fig. 8(b) shows the average error computed using the best 20 samples. This error is higher than that of the best sample. Note that this is not an obvious

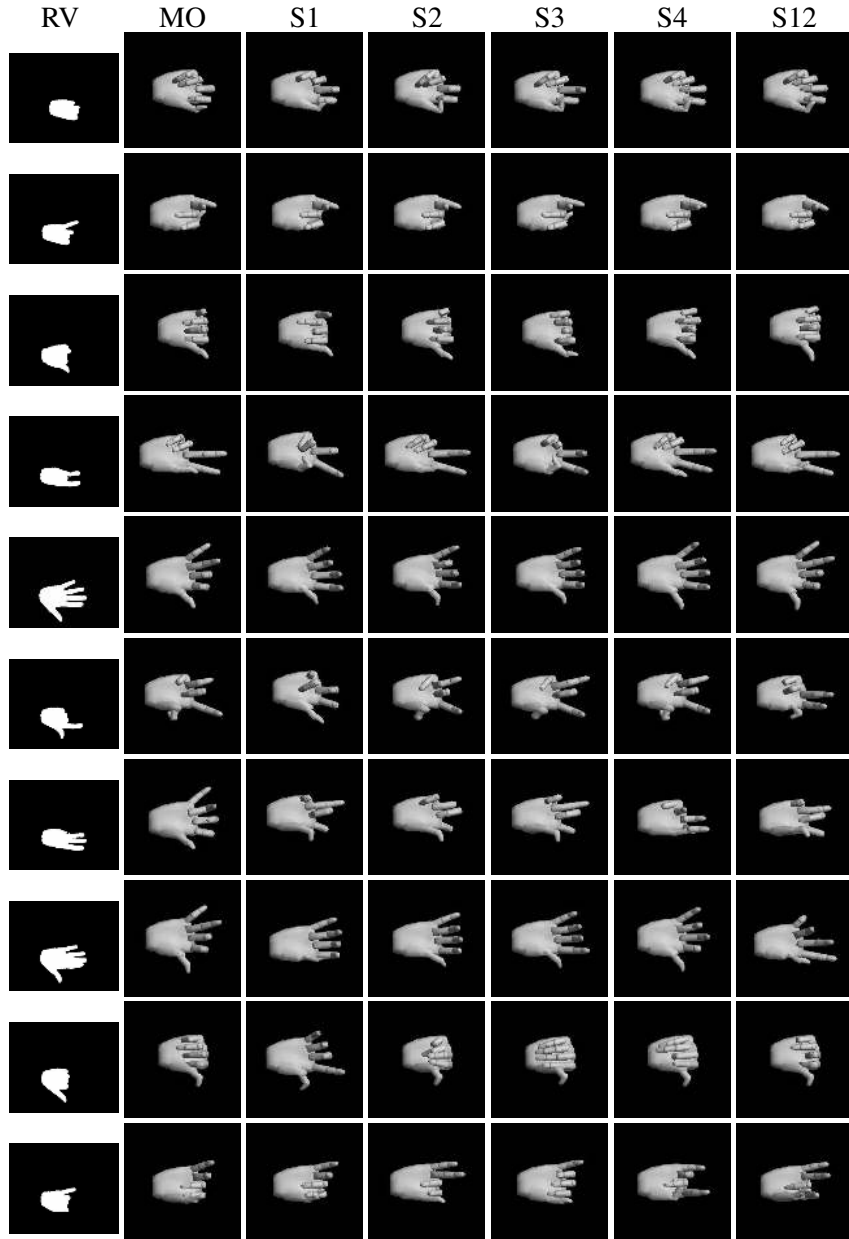


Figure 6: Example estimated hand poses obtained using the Multiple Sample (MS) approach using real video (RV). The feedback function was estimated from data.

result given that the best sample is determined without having knowledge of ground-truth. In fact, if the average error of the best 20 samples were lower than that of the best sample, then we could infer that our algorithm is very inaccurate at determining what samples are better. Thus this result positively endorses our MS algorithm.

For comparison, we used the ground-truth to select the best sample, based on minimum error. In other words, we have an oracle that picks the sample closest to the ground-truth. The resulting performance graph is shown in Fig. 8(c). This represents the lower-bound on the reconstruction error using the learned forward

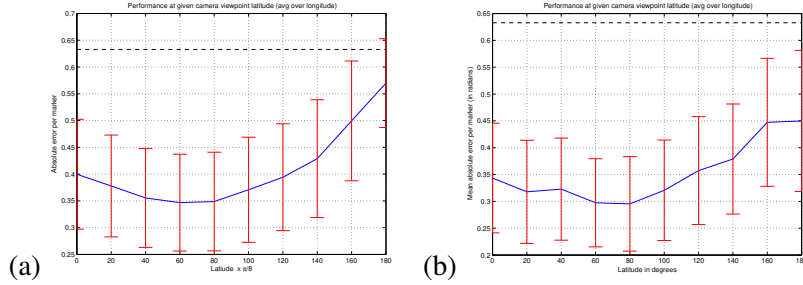


Figure 7: Mean Output (MO) inference performance for unrestricted view tests at given viewpoint latitudes (averaging over longitude). The feedback function is (a) the estimated $\hat{\zeta}$ (b) the computer graphics rendering ζ . A frontal view of the hand palm is at latitude $\beta_1 = \pi/2$, longitude $\beta_2 = 0$. For reference, the performance of an algorithm that chooses the estimate at random from the training data is shown. The angle range is in average 1.87 radians.




Figure 8: Multiple Samples (MS) inference for unrestricted view tests at given viewpoint latitudes (averaging over longitude). Feedback functions is the estimated $\hat{\zeta}$. A frontal view to the hand palm is at latitude $\beta_1 = \pi/2$, longitude $\beta_2 = 0$. (a) Most probable sample. (b) Average over all samples (20 most probable samples taken). (c) Best sample (determined using ground-truth information for comparison). For reference, the performance of an algorithm that chooses the estimate at random from the training data is shown. The angle range is in average 1.87 radians.

model. The graph is interesting in the sense that it separates the errors from the forward and feedback models.

7.2.2 Experiments with Real Images

We test our approach using video of hands (in any orientation) collected from a single uncalibrated camera. Pose estimates from 40 frames (taken every 0.9 secs apart) obtained via the MO approach are shown in Fig. 9. Note that there are incorrectly-segmented hands in this sequence. We decided to leave these in to avoid frame rearrangements (losing the uniform frame sampling), to show that segmentation does not always work correctly, and to show that this approach is inherently robust to extreme segmentation errors. In this experiment, there was usually visual agreement between reconstruction and estimate as seen in the figure. Note that even for a human observer, looking at the segmented silhouettes in the figure, reconstruction is sometimes ambiguous. There are also some configurations for which the system did not perform correctly.

Fig. 10 shows the estimates obtained via the MS approach. The frames shown were taken approximately

every 0.9 seconds. In the second row, we can see some limitations of the Hu moment feature space: sometimes, different hand orientations are very similar in the feature space. These apparently different hypotheses are close to each other in terms of their probability, given the features. The same effect repeats clearly in the third and sixth row. This problem might be alleviated by using a different input feature space. At an extreme one might consider the full silhouette as a feature. Of course there are important trade-offs to take into account when considering different features; e.g., invariants, and dimensionality.




Figure 9: 40 examples of estimated hand poses captured every 0.9 secs from real video (RV). Reconstruction found using the Mean Output (MO) approach. The feedback function was computed using computer graphics rendering.

7.3 2D Human Body Pose Reconstruction

In order to show that our approach can be employed, with no change, to perform other similar tasks (possibly with a different representation), here we now conduct performance tests in the task of estimating human body pose from a single image. The goal is to estimate the 2D locations of body markers in the image, given visual features computed from the person’s silhouette. In this experiment, we use the *Body-All-Views* dataset, which contains a total of over 100,000 samples. Of these, 8,000 were used for training and the rest for testing. The input-output pairs were defined as follows. The input consisted of the 10 Alt moments computed from the silhouette. The output consisted of 20 2D marker positions (40 DOF), which were then linearly encoded by nine values using PCA.




Figure 10: Example estimated hand poses obtained using the Multiple Sample (MS) approach and real video (RV). The feedback function was computed using computer graphics rendering.

The number of mixture components in the discriminative model was set to 15. This number was determined via the MDL criterion, exactly as before. Each function is a MLP with seven hidden nodes.

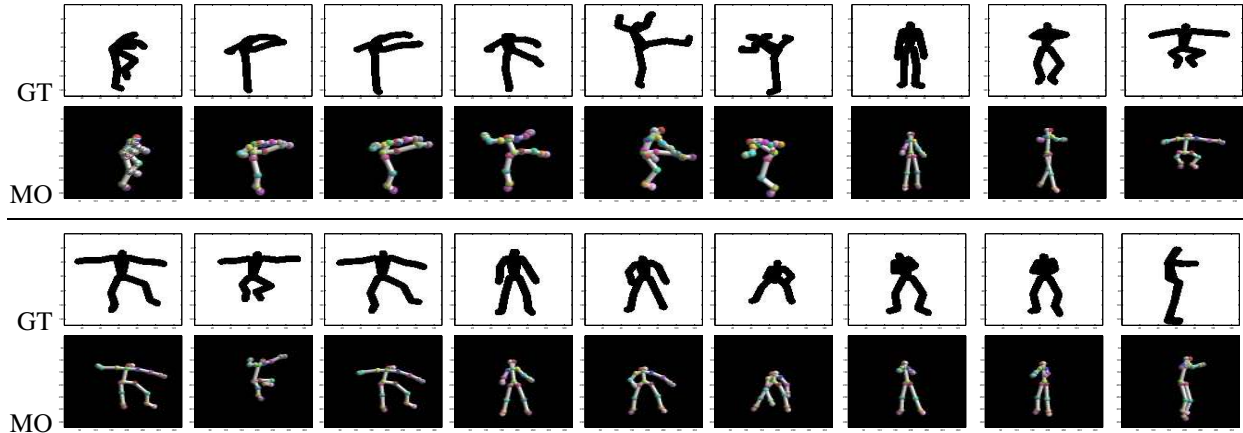


Figure 11: Example reconstruction of frames from test sequences with computer graphics-generated silhouettes.




Figure 12: Root mean-square-error (divided by number of markers) and variance per camera viewpoint (every $2\pi/32$ rads.). Units are percentage of body height. Approx. 110,000 test poses were used.

7.3.1 Quantitative Results

Fig. 11 shows the reconstruction obtained with the MO approach for frames taken from three synthetic sequences excluded from the training set. The agreement between reconstruction and observation is easy to perceive for all frames. Also, for self-occluding configurations, the estimate is still similar to ground-truth.

Fig. 12 shows the average marker error and variance per body orientation in percentage of body height. Note that the error is bigger for orientations closer to 0 and π radians. This intuitively agrees with the notion that at those angles (side-views), there is less visibility of the body parts. We consider this performance promising, given the complexity of the task and the simplicity of the approach. Just as a reference point, by choosing poses at random from those in the training set, the RMSE was 10.35% of body height (with a standard deviation of 4.4%). In related work, quantitative performance has usually been ignored, in part due to the lack of ground-truth and standard evaluation datasets.

7.3.2 Experiments with Real Images

We now test the approach using real video sequences of human body motion. We use the basic segmentation approach described in Sec. 6.2.3 to obtain silhouettes. Fig. 13 shows examples of system performance obtained via the MO approach for several relatively complex motion sequences. Even though the characteristics of the segmented body differ from the ones used for training, good performance is still achieved. Most reconstructions are visually close to what can be thought of as the right pose reconstruction. Body orientation is also accurate. In the Figure, we can see two particularly difficult configurations at the second row of real video (RV) images, fourth-sixth columns; the arm configuration is difficult to estimate. This could be due to the lack of relevant training data, as a consequence the discriminative model q may not approximate the generative model p very well around the input vector. In general, an important issue to keep in mind is that the visual differences between the rendered model and the real body observed could become critical and thus accurate rendering may be desirable. This varies from application to application; however in any case the general inference approach presented here remains the same.

In this work, we did not pursue use of a more realistic human body renderer. Due to differences in shape and width of body components observed in training versus testing, the visual features may differ. This is a relevant point since in almost all learning models, it is expected that the training data be a good approximation to the real test data. Improving the match between visual features used in training and testing, and thus potentially the overall performance, is an area that we plan to investigate in future research. Despite the fact that we have ignored differences in anthropometric characteristics between CG and real silhouettes, the performance observed for both articulated objects (hands - human bodies) is excellent given that only a single image is assumed available.




Figure 13: Reconstruction obtained from observing a human subject (every 10th frame).

8 Conclusions

In this paper, we have described a novel method that allows us to infer 3D and 2D articulated body pose from observed visual features in a single image, a problem usually regarded as ill-posed. This was done by combining generative and discriminative models to solve the complex probabilistic inference problem. This approach is most useful when the generative model is accurate (*e.g.*, we have an inverse mapping function) but it is difficult to perform inference using this model alone.

In order to solve the inference problem (and also perform MAP estimation), we have shown that a mathematically sound approach is to use a discriminative model and learn its parameters using relevant training data. The probability distribution implied by the discriminative model can be used as a proposal distribution to generate samples and find a posterior probability distribution (perform approximate inference) under the (accurate but complex) generative model.

When comparing it to other relevant methods, we can find alternative (dual) interpretations of this framework. The use of a generative model (through ζ) affords an alternative to complex discriminative models; for example, it is an alternative to the gating networks of the Mixture of Experts paradigm [23]. In general, instead of learning increasingly complex discriminative models such as [16, 12], we can exploit an accurate generative model and learn a simpler discriminative model. A clear advantage of using a generative model in this way is that it can provide useful information on the structure of the problem, a structure that discriminative models try to blindly uncover.

Our approach was demonstrated in a computer vision system that can estimate the articulated pose parameters of a human body or human hands, given features computed from a single image. This is a particularly difficult problem because this mapping is highly ambiguous, complex and it is infeasible to perform inference using the discriminative model. We have obtained promising results even using a very simple set of image features, such as moment invariants of the body silhouette. Choosing the best subset of image features for this application is by itself a complex problem, and a topic of ongoing research.

This approach offers several advantages over many previous methods for articulated pose estimation. These have tried in numerous ways to use camera geometry and/or model registration to perform pose estimation, resulting in iterative procedures that require careful choice of initial conditions (model placement). We have shown how in some cases these alternative approaches could be seen as inferring a posterior distribution using the generative model only. We have used camera geometry for defining our generative model but have not tried to solve the resulting optimization problem directly, instead we have had the help of the proposed discriminative model. In our approach no iterative minimization methods are used in pose infer-

ence. Moreover, inference is fully automatic – no manual initialization of the articulated model is required. Another set of previous approaches attempt to learn articulated model dynamics [5, 18, 40]; however, learning dynamics requires substantially more training data, and tends to produce systems that are biased towards specific motions. Our framework avoids this and infer pose from a single image only. However it is clear that in highly constrained environments, models of dynamics can provide an enormous advantage. In this paper we have chosen not to restrict the application domain yet, but we have proposed a general method that could be extended in different directions, one of them could involve taking into account the dynamics if the application domain thus allows.

Several interesting problems remain for future work. Within the context of articulated pose estimation. For example, (1) adapt the system to a specific body morphology, one of the major issues affecting performance and (2) integration of pose estimation with image segmentation for potentially greater robustness to occlusion and noise. Methods for incorporating knowledge of dynamics in the same framework should be investigated, as discussed in [33]. Another general problem is how to learn what the best (*e.g.*, visual) features are for specific problems or datasets. While promising advances have been made, extension of our framework to incorporate such concepts remains a topic for future investigation.

Acknowledgments

The hand sequences used in our experiments were collected in collaboration with Vassilis Athitsos at Boston University. We thank Tommi Jaakkola at MIT, Quaid Morris at University of Toronto, and Matt Brand at MERL for their valuable suggestions and interesting discussions. This research was supported in part by the U.S. Office of Naval Research under grants N000140310108 and N000140110444, and the U.S. National Science Foundation under grants IIS-0208876 and IIS-9809340.

References

- [1] F.L. Alt. Digital pattern recognition by moments. *Journal of the Association for Computing Machinery*, 9(2):240-258, April 1962.
- [2] S. I. Amari. Information geometry of the EM and *em* algorithms for neural networks. *Neural Networks*, 8(9):1379–1408, 1995.
- [3] C. Barron and I. Kakadiaris. Estimating anthropometry and pose from a single image. In *Proc. Computer Vision and Pattern Recognition*, pages 669–676, 2000.

- [4] M. Black and Y. Yacoob. Tracking and recognizing rigid and non-rigid facial motion using local parametric models of image motion. In *Proc. International Conference on Computer Vision*, 1995.
- [5] M. Brand. Shadow puppetry. In *Proc. International Conference on Computer Vision*, pages 1237–1244, 1999.
- [6] C. Bregler. Tracking people with twists and exponential maps. In *Proc. Computer Vision and Pattern Recognition*, pages 8–15, 1998.
- [7] J. Cheng and M. Druzdzel. AIS-BN: An adaptive importance sampling algorithm for evidential reasoning in large bayesian networks. *Journal of Artificial Intelligence Research*, 13:155–188, 2000.
- [8] I. Csiszar and G. Tusnady. Information geometry and alternating minimization procedures. *Statistics and Decisions*, 1:205–237, 1984.
- [9] A. Dempster, N. Laird, and D. Rubin. Maximum likelihood estimation from incomplete data. *Journal of the Royal Statistical Society (B)*, 39(1):1–38, 1977.
- [10] J. Deutscher, A. Blake, and I. Reid. Articulated body motion capture by annealed particle filtering. In *Proc. Computer Vision and Pattern Recognition*, 2000.
- [11] P. Felzenszwalb and D. Huttenlocher. Efficient matching of pictorial structures. In *Proc. Computer Vision and Pattern Recognition*, 2000.
- [12] J.H. Friedman. Multivariate adaptive regression splines. *The Annals of Statistics*, 19,1-141, 1991.
- [13] D. Gavrilu and L. Davis. Towards 3-d model-based tracking and recognition of human movement: a multi-view approach. In *International Workshop on Automatic Face and Gesture Recognition*, pages 272–277, 1995.
- [14] I. Haritaoglu, D. Harwood, and L. Davis. Ghost: A human body part labeling system using silhouettes. In *International Conference on Pattern Recognition*, pages 77–82, 1998.
- [15] T. Heap and D. Hogg. Towards 3d hand tracking using a deformable model. In *Proc. International Conference on Automatic Face and Gesture Recognition*, pages 140–145, 1996.
- [16] G. Hinton, B. Sallans, and Z. Ghahramani. A hierarchical community of experts. *Learning in Graphical Models*, M. Jordan (editor), pages 479–494, 1998.

- [17] D. Hogg, S. Dudani, K. Breeding, and R. McGhee. Model-based vision: A program to see a walking person. *Image and Vision Computing*, 1(1):5–20, 1983.
- [18] N.R. Howe, M.E. Leventon, and W.T. Freeman. Bayesian reconstruction of 3d human motion from single-camera video. In *Advances in Neural Information Processing Systems*, volume 12, pages 820–826, 2000.
- [19] M. K. Hu. Visual pattern recognition by moment invariants. *IRE Transactions Information Theory*, IT(8):179–187, 1962.
- [20] M. Isard and A. Blake. Condensation – conditional density propagation for visual tracking. *International Journal of Computer Vision*, 29(1): 5-28, 1998.
- [21] G. Johansson. Visual perception of biological motion and a model for its analysis. *Perception and Psychophysics*, 14(2): 210-211, 1973.
- [22] M. Jordan. *Learning in graphical models*. Kluwer Academic, The Netherlands, 1999.
- [23] M. Jordan and R. Jacobs. Hierarchical mixtures of experts and the EM algorithm. *Neural Computation*, 6, 181-214, 1994.
- [24] D. Mackay. Introduction to Monte Carlo methods. *Learning in Graphical Models*, 1998.
- [25] G. J. McLachlan. *Discriminant Analysis and Statistical Pattern Recognition*. Wiley, New York, 1992.
- [26] R. Neal and G. Hinton. A view of the em algorithm that justifies incremental, sparse, and other variants. *Learning in Graphical Models*, M. Jordan (editor), pages 355–368, 1998.
- [27] A. Ng and M. Jordan. On discriminative vs. generative classifiers: a comparison of logistic regression and naive Bayes. In *Advances in Neural Information Processing Systems*, pages 841–848, 2001.
- [28] D. Ormoneit, H. Sidenbladh, M. Black, and T. Hastie. Learning and tracking cyclic human motion. In *Advances in Neural Information Processing Systems 13*, pages 894–900, 2001.
- [29] V. Pavlović, J. Rehg, and J. MacCormick. Learning switching linear models of human motion. In *Advances in Neural Information Processing Systems 13*, pages 981–987, 2001.
- [30] J. Pearl. *Probabilistic Reasoning in Intelligent Systems*. Morgan-Kaufman, 1988.
- [31] J. M. Rehg and T. Kanade. Model-based tracking of self-occluding articulated objects. In *Proc. International Conference on Computer Vision*, pages 612–617, 1995.

- [32] J. Rissanen. Stochastic complexity and modeling. *Annals of Statistics*, 14,1080-1100, 1986.
- [33] R. Rosales. *The Specialized Mappings Architecture, with Applications to Vision-Based Estimation of Articulated Body Pose*. PhD thesis, Boston University, 2002.
- [34] R. Rosales, V. Athitsos, L. Sigal, and S. Sclaroff. 3d hand pose estimation using specialized mappings. In *Proc. International Conference on Computer Vision*, pages 378–387, 2001.
- [35] R. Rubinstein. *Simulation and the Monte Carlo method*. John Wiley & Sons, 1981.
- [36] Y. Rubinstein and T. Hastie. Discriminative vs. informative learning. In *3rd International Conference on Knowledge Discovery and Data Mining*, pages 49–56, 1997.
- [37] N. Shimada, Y. Shirai, Y. Kuno, and J. Miura. Hand gesture estimation and model refinement using monocular camera - ambiguity limitation by inequality constraints. In *Proc. International Conference on Automatic Face and Gesture Recognition*, pages 268–273, 1998.
- [38] L. Sigal, S. Sclaroff, and V. Athitsos. Estimation and prediction of evolving color distributions for skin segmentation under varying illumination. In *Proc. Computer Vision and Pattern Recognition*, pages 152–159, 2000.
- [39] C. Sminchisescu and B. Triggs. Covariance scaled sampling for monocular 3d body tracking. In *Proc. Computer Vision and Pattern Recognition*, pages 447–454, 2001.
- [40] Y. Song, X. Feng, and P. Perona. Towards detection of human motion. In *Proc. Computer Vision and Pattern Recognition*, pages 810–817, 2000.
- [41] C. J. Taylor. Reconstruction of articulated objects from point correspondences in a single uncalibrated image. *Computer Vision and Image Understanding: CVIU*, 80(3):349–363, December 2000.
- [42] Virtual Technologies, Inc., Palo Alto, CA. *VirtualHand Software Library Reference Manual*, 1998.
- [43] C. Wren, A. Azarbayejani, T. Darrell, and A. Pentland. Pfinder: Real time tracking of the human body. *PAMI*, 19(7):780-785, 1997.
- [44] S.C. Zhu, C. Guo, and Y. Wu. Modeling visual patterns by integrating descriptive and generative models. *International Journal of Computer Vision*, 53(1):5–29, 2003.