THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Progressive search space reduction for human pose estimation

Citation for published version:

Ferrari, V, Marin-Jimenez, M & Zisserman, A 2008, Progressive search space reduction for human pose estimation. in *Computer Vision and Pattern Recognition, 2008. CVPR 2008. IEEE Conference on*. Institute of Electrical and Electronics Engineers (IEEE), pp. 1-8. <https://doi.org/10.1109/CVPR.2008.4587468>

Digital Object Identifier (DOI):

[10.1109/CVPR.2008.4587468](https://doi.org/10.1109/CVPR.2008.4587468)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

Published In:

Computer Vision and Pattern Recognition, 2008. CVPR 2008. IEEE Conference on

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

Progressive Search Space Reduction for Human Pose Estimation

Vittorio Ferrari
University of Oxford
ferrari@robots.ox.ac.uk

Manuel Marín-Jiménez
University of Granada
mjmarin@decsai.ugr.es

Andrew Zisserman
University of Oxford
az@robots.ox.ac.uk

Abstract

The objective of this paper is to estimate 2D human pose as a spatial configuration of body parts in TV and movie video shots. Such video material is uncontrolled and extremely challenging.

We propose an approach that progressively reduces the search space for body parts, to greatly improve the chances that pose estimation will succeed. This involves two contributions: (i) a generic detector using a weak model of pose to substantially reduce the full pose search space; and (ii) employing ‘grabcut’ initialized on detected regions proposed by the weak model, to further prune the search space. Moreover, we also propose (iii) an integrated spatio-temporal model covering multiple frames to refine pose estimates from individual frames, with inference using belief propagation.

The method is fully automatic and self-initializing, and explains the spatio-temporal volume covered by a person moving in a shot, by soft-labeling every pixel as belonging to a particular body part or to the background. We demonstrate upper-body pose estimation by an extensive evaluation over 70000 frames from four episodes of the TV series Buffy the vampire slayer, and present an application to full-body action recognition on the Weizmann dataset.

1. Introduction

Our aim is to detect and estimate 2D human pose in video, i.e. recover a distribution over the spatial configuration of body parts in every frame of a shot. Various pose representations can then be derived, such as a soft-labelling of every pixel as belonging to a particular body part or the background (figure 1b); or the ‘stickman’ of figure 1c, indicating the location, orientation, and size of body parts. Note, our objective here is not to estimate 3D human pose as in [4, 23, 26].

We wish to obtain pose estimates in highly challenging uncontrolled imaging conditions, typical of movies and TV shows (figure 7). Achieving this is one of the main contributions of the paper. In this setting, images are often very cluttered, and a person might cover only a small proportion of the image area, as they can appear at any scale. Illumi-

Figure 1. **Objective of this work.** (a) Input image. (b) Soft-labelling of pixels to body parts or background. Red indicates torso, green upper arms, blue lower arms and head. Brighter pixels are more likely to belong to a part. Color planes are added up, so that purple indicates overlap between lower-arm and torso, yellow between upper-arm and torso, and so on. (c) Stickman representation of pose, obtained by fitting straight line segments to the segmentations in (b). For enhanced visibility, the lower arms are in yellow and the head is in purple.

nation varies over a diverse palette of lighting conditions, and is often quite dark, resulting in poor image contrast. A person’s appearance is unconstrained, as they can wear any kind of clothing, including body-tight or loose, short or long sleeves, and any colors/textures. The background is unknown and changes over time, preventing the use of background subtraction techniques [3, 6]. Finally, the camera is usually moving, causing motion blur, and multiple people can be present at the same time and can occlude each other during a shot.

Our method covers all poses within the upper-body frontal range. Special attention is given to the arms, as they carry most of the information necessary to distinguish pose. The proposed method supports arms folded over the torso, stretching outwards, pointing forward, etc.

The need for such human centered tracking is evident, with applications ranging from video understanding and search through to surveillance. Indeed 2D human segmentation is often the first step in determining 3D human pose from individual frames [2]. We illustrate the use of the extracted poses with an application to action recognition on the Weizmann dataset.

1.1. Approach overview

We overview the method here for the upper-body case, where there are 6 parts: head, torso, and upper/lower right/left arms (figure 1). Full details are given in section 2.

The method is also applicable to full bodies, as demonstrated in section 4.

A recent and successful approach to 2D human tracking in video has been to detect in every frame, so that tracking reduces to associating the detections [19, 25]. We adopt this approach where detection in each frame proceeds in three stages, followed by a final stage of transfer and integration of models across frames.

In our case, the task of pose detection is to estimate the parameters of a 2D articulated body model. These parameters are the (x, y) location of each body part, its orientation θ , and its scale. Assuming a single scale factor for the whole person, shared by all body parts, the search space has $6 \times 3 + 1 = 19$ dimensions. Even after taking into account kinematic constraints (e.g. the head must be connected to the torso), there are still a huge number of possible configurations.

Since at the beginning we know nothing about the person’s pose, clothing appearance, location and scale in the image, directly searching the whole space is a time consuming and very fragile operation (there are too many image patches that could be an arm or a torso!). Therefore, in our approach the first two stages use a weak model of a person obtained through an upper-body detector generic over pose and appearance. This weak model only determines the approximate location and scale of the person, and roughly where the torso and head should lie. However, it knows nothing about the arms, and therefore very little about pose. The purpose of the weak model is to *progressively reduce the search space* for body parts.

The next two stages then switch to a stronger model, i.e. a pictorial structure [6, 18, 19] describing the spatial configuration of all body parts and their appearance. In the reduced search space, this stronger model has much better chances of inferring detailed body part positions.

1. Human detection. We start by detecting human upper-bodies in every frame, using a sliding window detection based on Histograms of Oriented Gradients [5], and associate detections over time. Each resulting track carves out of the total spatio-temporal volume the smaller subvolume covered by a person moving in the shot. This reduces the search space by by setting bounds on the possible (x, y) locations of the body parts and by fixing their scale, thus removing a dimension of the search space entirely.

2. Foreground highlighting. At this stage the search for body parts is only limited by the maximum extent possible for a human of that scale centered on the detected position. We restrict the search area further by exploiting prior knowledge about the structure of the detection window. Relative to it, some areas are very likely to contain part of the person, whereas other areas are very unlikely. This allows the initialization of a Grabcut segmentation [20], which removes part of the background clutter. This stage

further constrains the search space by limiting the (x, y) locations to lie within the area output by Grabcut.

3. Single-frame parsing. We obtain a first pose estimate based on the *image parsing* technique of Ramanan [18]. The area to be parsed is restricted to the region output of foreground highlighting. Since the person’s scale has been fixed by stage 1), no explicit search for body parts over scales is necessary.

4. Spatio-temporal parsing. The appearance of the body parts of a person changes little within a shot. Moreover, the position of body parts changes smoothly over time. We exploit both kinds of temporal continuity in a second pose estimation procedure which (i) uses appearance models integrated from multiple frames where the system is confident about the estimated pose; and (ii) infers over a joint spatio-temporal model of pose, capturing both kinematic constraints within a frame, and temporal continuity constraints between frames. As appearance is a powerful cue about the location of parts, the better appearance models improve results for frames where parsing failed or is inaccurate. At the same time, the spatio-temporal model tightens the posterior distributions over part positions and disambiguates multiple modes hard to resolve based on individual frames.

1.2. Related works

Our work builds mainly on the *Learning to Parse* approach by Ramanan *et al.* [18], which provides the pictorial structure inference engine we use in stage 3, and on the *Strike-a-pose* work [19]. The crucial difference to both works is the way the search space of body part configurations is treated. Thanks to the proposed detection and foreground highlighting stages, we avoid the very expensive and fragile search necessary in [18, 19]. Moreover, compared to [19], our initial detection stage is *generic over pose*, so we are not limited to cases where the video contains a pre-defined characteristic pose at a specific scale. We also generalize and improve the idea of transferring appearance models of [19]. Rather than using a single frame containing the characteristic pose, we integrate models over multiple frames containing any pose.

Initially, our approach attempts to fit a pictorial structure model independently to each frame (stage 3). Previous use of these models have tolerated only limited amounts of background clutter [6, 18] and often assume knowledge of the person’s scale [18, 19] or background subtraction [6]. A few methods operate interactively from regions of interest provided by the user [14].

There are also methods that detect humans using generative models for the entire video sequences, e.g. [9, 11]. However, todate these methods have been limited to relatively simple backgrounds and to no occlusion of the person.

Our spatio-temporal model (section 2.4) is most closely related that of [23, 24], but our framework is fully automatic (it does not need any manual initialization or background subtraction).

Finally, [15] recovers unusual, challenging body configurations in sports images by combining segmentation and detectors trained to specific body parts, but requires a person centered in the image and occupying most of it.

2. Technical details

2.1. Upper-body detection and temporal association

Upper-body detection. In most shots of movies and TV shows, only the upper-body is visible. To cope with this situation, we have trained an upper-body detector using the approach of Dalal and Triggs [5], which achieves excellent performance on the related task of full-body pedestrian detection. Image windows are spatially subdivided into tiles and each is described by a Histogram of Oriented Gradients. A sliding-window mechanism then localizes the objects. At each location and scale the window is classified by an SVM as containing the object or not. Photometric normalization within multiple overlapping blocks of tiles makes the method particularly robust to lighting variations.

The training data consists of 96 video frames from three movies (*Run Lola run*, *Pretty woman*, *Groundhog day*), manually annotated with a bounding-box enclosing an upper-body. The images have been selected to maximize diversity, and include many different actors, with only a few images of each, wearing different clothes and/or in different poses. No images from the test material (shots from *Buffy the Vampire Slayer*) were used for training.

Following Laptev [12], the training set is augmented by perturbing the original examples with small rotations and shears, and by mirroring them horizontally. This improves the generalization ability of the classifier. By presenting it during training with misalignments and variations, it has a better chance of noticing true characteristics of the pattern, as opposed to details specific to individual images. The augmented training set is 12 times larger and contains more than 1000 examples.

We choose an operating point of 90% detection-rate at 0.5 false-positives per image. This per-frame detection-rate translates into an almost perfect per-track detection-rate after temporal association (see below). Although individual detections might be missed, entire tracks are much more robust. Moreover, we remove most false-positives by weeding out tracks shorter than 20 frames.

In practice, this detector works well for viewpoints up to 30 degrees away from straight frontal, and also detects back views (figure 7).

Temporal association. After applying the upper-body detector to every frame in the shot independently, we associate

Figure 2. **Overview of the single-frame steps.** **1. Upper body detection:** The detected person (inner rectangle) and enlarged window where further processing is applied (outer rectangle). **2. Foreground highlighting:** (b) subregions for initializing Grabcut. (c) foreground region output by Grabcut. **3. Parsing:** (d) area to be parsed $\mathcal{F}$ (dilated from (c)) and (e) edges within $\mathcal{F}$. (f) posterior of the part positions $p(l_i|I)$ after the edge-based inference. (g) posterior after the second inference, based on edges and appearance.

the resulting bounding-boxes over time by maximizing their temporal continuity. This produces *tracks*, each connecting detections of the same person.

Temporal association is cast as a grouping problem [25], where the elements to be grouped are bounding-boxes. As similarity measure we use the area of the intersection divided by the area of the union (IoU), which subsumes both location and scale information, damped over time. We group detections based on these similarities using the Clique Partitioning algorithm of [7], under the constraint that no two detections from the same frame can be grouped. Essentially, this forms groups maximizing the IoU between nearby time frames.

This algorithm is very rapid, taking less than a second per shot, and is robust to missed detections, because a high IoU attracts bounding-boxes even across a gap of several frames. Moreover, the procedure allows persons to overlap partially or to pass in front of each other, because IoU injects a preference for *continuity scale* in the grouping process, in addition to location, which acts as a disambiguation factor.

In general, the ‘detect & associate’ paradigm is substantially more robust than regular tracking, as recently demonstrated by several authors [17, 25].

2.2. Foreground highlighting

The location and scale information delivered by an upper-body detection greatly constrains the space of possible body parts. They are now confined to the image area surrounding the detection, and their approximate size is known, as proportional to the detection’s scale. However, to accommodate for all possible arm poses we must still explore a sizeable area (figure 2a). Stretching out the arms

in any direction forms a large circle centered between the shoulders. In challenging images from TV shows, this area can be highly cluttered, confusing the body part estimator.

Fortunately, we have *prior knowledge* about the structure of the search area. The head lies somewhere in the middle upper-half of the detection window, and the torso is directly underneath it (figure 2b). In contrast the arms could be anywhere. This is known because the detector has been explicitly trained to respond to such structures. We propose to exploit this knowledge to initialize Grabcut [20], by learning initial foreground/background color models from regions where the person is likely to be present/absent. The resulting segmentation removes much of the background clutter, substantially simplifying the later search for body parts (figure 2c).

Let $\mathcal{R}$ be a region of interest obtained by enlarging the detection window as in figure 2a. $\mathcal{R}$ is divided into four subregions F, F_c, B, U (see figure 2b). Grabcut is initialized as follows: the foreground model is learnt from F and F_c (F_c is known to belong to the person, while F contains mostly foreground, but some background as well); and the background model from B (it covers mostly background, but it might also include part of the arms, depending on the pose). Furthermore, the region F_c is clamped as foreground, but grabcut is free to set pixel labels in all other subregions (we have extended the original Grabcut algorithm to enable these operations). The U region is neutral and no color model is learnt from it. The setup accurately expresses our prior knowledge and results in a controlled, upper-body-specific segmentation, assisted by as much information as we can derive from the previous object detection process. Near the head, B and F_c compete for the U region, with the foreground growing outwards until it meets a background-colored area, resulting in a good head segmentation. Along the sides, the background floods into the initial F to segment the shoulders, while at the same time the arms get labeled as foreground because they are colored more similarly to the initial F than to the initial B (figure 3).

The above procedure is rather conservative, and it often retains parts of the background. The goal is not to achieve a perfect segmentation, but to reduce the amount of background clutter (figure 3). It is more important not to lose body parts, as they cannot be recovered later. To validate this, we have inspected 1584 frames of a *Buffy* episode (i.e. every 10th frame) and only in 71 a body part was lost (4.5%). In contrast to traditional background subtraction, used in many previous works to extract silhouettes [3, 6, 8], our method does not need to know the background *a priori*, and allows the background to change over time (in video).

2.3. Single-frame parsing

Our main goal is to explain the spatio-temporal volume covered by a person moving in a shot. In particular, we

Figure 3. Examples of foreground highlighting.

want to estimate the 2D pose of the person, as the location, orientation and size of each body part. Ideally, the exact image regions covered by the parts should also be found. For estimating 2D pose in individual video frames, we build on the *image parsing* technique of Ramanan [18]. In the following we first briefly summarize it, and then describe our extensions.

Image parsing [18]. A person is represented as a pictorial structure composed of body parts tied together in a tree-structured conditional random field. Parts, l_i , are oriented patches of fixed size, and their position is parametrized by location and orientation. The posterior of a configuration of parts $L = \{l_i\}$ given an image I can be written as a log-linear model

$$P(L|I) \propto \exp \left(\sum_{(i,j) \in E} \Psi(l_i, l_j) + \sum_i \Phi(l_i) \right) \quad (1)$$

The binary potential $\Psi(l_i, l_j)$ corresponds to a spatial prior on the relative position of parts and embeds the kinematic constraints (e.g. the upper arms must be attached to the torso). The unary potential $\Phi(l_i)$ corresponds to the local image evidence for a part in a particular position (likelihood). Since the model structure E is a tree, inference is performed exactly and efficiently by sum-product Belief Propagation.

The key idea of [18] lies in the special treatment of Φ . Since the appearance of neither the parts nor the background is known at the start, only edge features are used. A first inference based on edges delivers soft estimates of body part positions, which are used to build appearance models of the parts (e.g. in figure 2f the torso is in red). Inference is then repeated using an updated Φ incorporating both edges and appearance. The process can be iterated further, but in this paper we stop at this point. The technique is applicable to quite complex images because (i) the appearance of body parts is a powerful cue, and (ii) appearance models can be learnt from the image itself through the above two-step process.

The appearance models used in [18] are color histograms over the RGB cube discretized into $16 \times 16 \times 16$ bins. We refer to each bin as a *color* c . Each part l_i has foreground and background likelihoods $p(c|fg)$ and $p(c|bg)$. These are learnt from a part-specific soft-assignment of pixels to foreground/background derived from the posterior of the part position $p(l_i|I)$ returned by parsing. The posterior for a

Figure 4. The *image parsing* pose estimation algorithm of [18] applied to the image in figure 1. (a) All edges inside region $\mathcal{R}$, without filtering them through foreground highlighting. (b) Parsing applied to the whole region $\mathcal{R}$. It achieves a worse estimate than when helped by foreground highlighting, figure 2g, because it is attracted by the bars in the background. (c) Parsing applied directly to the whole image, without reducing the search space to $\mathcal{R}$ based on the initial person detection. It fails entirely.

pixel to be foreground given its color $p(fg|c)$ is computed using Bayes’ rule and used during the next parse.

As in [18], in our implementation we explicitly maintain a 3D binned volume to represent the possible (x, y, θ) positions of each part (discretization: every pixel for (x, y) and 24 steps for θ). This dense representation avoids the sampling needed by particle representations (e.g. [23, 24]).

When run unaided on a highly cluttered image such as figure 1a, without any idea of where the person might be or how large it is, parsing fails entirely (figure 4c). There are simply too many local image structures which could be a limb, a head, or a torso. This is assessed quantitatively in section 3.

We reduce the space explored by parsing based on three sources of information: (i) the *location and scale information* supplied by the upper-body detector, is used to define the enlarged search region $\mathcal{R}$. Parsing is run only within $\mathcal{R}$, rescaled to a fixed size, tuned to roughly yield the part sizes expected by the parser. Thanks to the proper use of scale information from detection, we effectively obtain scale-invariant pose estimation, without having to explicitly search for body parts at multiple scales. This significantly reduces ambiguity and false positive detections. (ii) *Foreground highlighting*. We further simplify pose estimation by restricting the area to be parsed to the region $\mathcal{F} \subset \mathcal{R}$ output of foreground highlighting (figure 2d). This is realized by removing all edges outside $\mathcal{F}$ and setting all pixels $\mathcal{R} \setminus \mathcal{F}$ to black. This causes the image evidence $\Phi(l_i)$ to go to $-\infty$ for $l_i \notin \mathcal{F}$, and hence it is equivalent to constraining the space of possible poses. (iii) *Head and torso constraints*. A final assistance is given by mildly constraining the location of the head and torso based on our prior knowledge about the spatial structure of $\mathcal{R}$ (see section 2.2). The constraints come in the form of broad subregions $\mathcal{H}, \mathcal{T} \in \mathcal{R}$ where the head and torso must lie, and are realized by setting $\Phi(l_{head}), \Phi(l_{torso})$ to $-\infty$ for $l_i \notin \mathcal{H}, \mathcal{T}$ (figure 2d). These constraints directly reflect our prior knowledge from the detection process and therefore do not limit the range of poses covered by the parser (e.g. for the arms).

Figure 5. **Spatio-temporal model.** For clarity, only head (l_H), torso (l_T), and left/right upper arms (l_{LA}, l_{RA}) are drawn.

All the above aids to pose estimation are made possible from the initial generic upper-body detection. Foreground highlighting and location constraints can only be automated when building on a detection window. The combined effect of these improvements is a vastly more powerful parser, capable of estimating 2D pose in a highly cluttered image, even when the person occupies only a small portion of it. Moreover, parsing is now faster, as it searches only an image region, supports persons at multiple scales, and multiple persons at the same time, as each detection is parsed separately.

2.4. Spatio-temporal parsing

Parsing treats each frame independently, ignoring the temporal dimension of video. However, all detections in a track cover the same person, and people wear the same clothes throughout a shot. As a consequence, the appearance of body parts is quite stable over a track. In addition to this continuity of appearance, video offers also continuity of geometry: the position of body parts changes smoothly between subsequent frames.

In this section, we exploit the continuity of appearance for improving pose estimations in particularly difficult frames, and the continuity of geometry for disambiguating multiple modes in the positions of body parts, which are hard to resolve based on individual frames.

Learning integrated appearance models. The idea is to find the subset of frames where the system is confident of having found the correct pose, integrate their appearance models, and use them to parse the whole track again. This improves pose estimation in frames where parsing has either failed or is inaccurate, because appearance is a strong cue about the location of parts.

Frames where parsing infers a highly confident configuration of body parts provide good reference appearance models (figure 6a). The measure of confidence used here is the entropy of the posterior of the part positions $p(l_i|I)$, accumulated over all parts to give the *total pose entropy* $TPE = -\sum_i H(p(l_i|I))$. Rather than simply selecting the single frame with the lowest TPE, we learn models by *integrating* over all frames with a similar low TPE. It can be shown [21] that the distribution minimizing the total KL-divergence to a set of distributions is their average.

Hence, we integrate the foreground and background likelihoods $\{p_r(c|fg)\}, \{p_r(c|bg)\}$ from the reference frames r by averaging them. The integrated posteriors $p_i(fg|c)$ are then obtained by applying Bayes' rule.

The integrated models are richer, in that $p_i(fg|c)$ is nonzero for a broader range of colors, so they generalize to a larger number of frames. Moreover, they are more accurate, because estimated over a wider support.

Spatio-temporal inference. We extend the single-frame person model of [18] to include dependencies between body parts over time (figure 5). The extended model has a node for every body part in every frame of a continuous temporal window (11 frames in our experiments). The posterior of all configurations of parts $\{L^t\} = \{l_i^t\}$ given all frames $\{I^t\}$ can be written as

$$P(\{L^t\}|\{I^t\}) \propto \exp \left(\sum_{t,i} \left(\sum_{j|(i,j) \in E} \Psi(l_i^t, l_j^t) + \Phi(l_i^t) + \Omega(l_i^t, l_i^{t+1}) \right) \right) \quad (2)$$

In addition to the *kinematic* dependencies Ψ between different parts in a single frame, there are *temporal* dependencies Ω between nodes representing the same part in subsequent frames. As a temporal prior $\Omega(l_i^t, l_i^{t+1})$ we use a simple box-shaped distribution limiting the difference in the $l_i^t = (x, y, \theta)$ position of a body part between frames. We use the integrated appearance models to obtain a better image likelihood Φ . Since the spatio-temporal graphical model has loops, we perform approximate inference with sum-product Loopy Belief Propagation, which in practice delivers a good estimate of the posterior marginals and is computationally efficient.

The spatio-temporal inference is a batch process treating all frames in the temporal window simultaneously, as opposed to traditional tracking, where only past frames can influence estimation in the current frame. The inference procedure outputs the full marginal posteriors $p(l_i^t|\{I^t\})$, defining the probability of every possible (x, y, θ) body part position in every frame. This is better than a single MAP solution [19], as any remaining ambiguity and uncertainty is visible in the full posteriors (e.g. due to blurry images, or tubular background structures colored like the person's arms). Finally, our joint spatio-temporal inference is better than simply smoothing the single-frame posteriors over time, as the kinematic dependencies within a frame and temporal dependencies between frames *simultaneously help each other*.

As a combined result of using integrated, high-confidence appearance models within a joint spatio-temporal inference, the final pose estimates are tighter and more accurate than the single-frame ones (figure 6, and compare figure 2g to figure 1b). Moreover, they are temporally continuous, which is useful for estimating the motion of body parts for action recognition.

Figure 6. **Spatio-temporal parsing.** (a) one of several frames with low TPE after single-frame parsing, from which integrated appearance models are learnt. (b) a frame with high TPE. The system is uncertain whether the right arm lies on the window or at its actual position. (c) Output of spatio-temporal parsing. Thanks to the integrated appearance models and the spatio-temporal inference, the right arm ambiguity is resolved, and the estimated position of the torso improves visibly.

3. Upper-body pose estimation results

We have applied our pose estimation technique to episodes 2,4,5 and 6 of season five of *Buffy the vampire slayer*, for a total of more than 70000 video frames over about 1000 shots.

The examples in figure 7, show that the proposed method meets the challenges set in the introduction. It successfully recovers the configuration of body parts in spite of extensive clutter, persons of different size, dark lighting and low contrast (b4, c3, f2). Moreover, the persons wear all kinds of clothing, e.g. ranging from long sleeves to sleeveless (b3, a3, a4), and this is achieved despite the fact that their appearance is unknown *a priori* and was reconstructed by the algorithm. The system correctly estimated a wide variety of poses, including arms folded over the body (a2, c1, c2), stretched out (c3, e3), and at rest (c4, e1). The viewpoint doesn't have to be exactly frontal, as the system tolerates up to about 30 degrees of out-of-plane rotation (c1). Persons seen from the back are also covered, as the upper-body detector finds them and we don't rely on skin-color segmentation (e1, f1). Finally, the method properly deals with multiple persons in the same image and delivers a separate pose estimate for each (e2, f4).

We quantitatively assess these results on 69 shots divided equally among three of the episodes¹. We have annotated the ground-truth pose for four frames spread roughly evenly throughout the shot, by marking each body part by one line segment [15] (figure 7a). Frames were picked where the person is visible at least to the waist and the arms fit inside the image. This was the sole selection criterion. In terms of imaging conditions, shots of all degrees of difficulty have been included. A body part returned by the algorithm is

¹The data is available on our website: www.robots.ox.ac.uk/~vgg/research/pose_estimation/index.html

Figure 7. **Pose estimation results.** (a1) example ground-truth ‘stickman’ annotation. All other subfigures are the output of the proposed method, with body part segmentations overlaid. For illustration, in (a2) we also overlay the stickman derived by our method. The color coding is as follows: head = purple, torso = red, upper arms = green, lower arms = yellow. A fair sample of failures are also included, e.g. (f1) is missed as a detection, and the wrong pose is obtained in (e3) (rear person). Notice how the leftmost person in (e2) is largely occluded.

considered correct if its segment endpoints lie within 50% of the length of the ground-truth segment from their annotated location. The initial detector found an upper-body in 88% of the $69 \times 4 = 276$ annotated frames. Our method correctly estimates 56% of the $243 \times 6 = 1458$ body parts in these frames.

To put this result in perspective, we have repeated the experiment with stripped down versions of the system. Pars-

ing without foreground highlighting, only exploiting the location and scale from the detection, causes performance to drop to 41%. We also compare to the image parser of [18] using software supplied by the author, and run unaided, directly on the image. [18] achieves only 10%, thus highlighting the great challenge posed by this data, and the substantial improvements brought by our technique. The results also confirm that both search space reduction stages we pro-

posed (starting from a detection and foreground highlighting) contribute considerably to the quality of the results.

4. Application: Action recognition on the Weizmann dataset

Determining human pose is often a first step to action recognition. For example [13] classifies frames by their pose as a pre-filter to spatio-temporal action detection, and [22] specifies actions by pose queries.

In this section we apply the extracted pose representation to the task of action recognition on the Weizmann dataset [3], which includes nine actions performed by nine different subjects who are fully visible in all frames (including their legs). Following the standard leave-one-out evaluation protocol [3, 8, 16], we train on eight subjects and test on the remaining one. The experiment is then repeated by changing the test subject and recognition rates are averaged.

Here, we replace the upper-body detector by the standard HOG based pedestrian detector of Dalal and Triggs [5], and employ a full-body pictorial structure including also upper and lower legs. Our action descriptor is obtained by subtracting soft-segmentations between pairs of subsequent frames, and accumulating the differences over the whole sequence. We train a linear SVM on this descriptor and use it to classify the actions of the test subject. Although previous works using background subtraction achieve perfect results on this dataset [3, 8], the only work we are aware of tackling the task without any static background assumption only obtains 73% recognition rate [16]. While operating in the same conditions, our method achieves the significantly higher rate of 88%. These results demonstrate the suitability of our technique to full body pose estimation.

5. Appraisal and future work

We have demonstrated automated upper body pose estimation on extremely challenging video material – the principal objective of this work.

The numerous works defining action descriptors based on body outlines [3, 8] could benefit from our technique, as it provides outlines without resorting to traditional background segmentation, requiring a known and static background.

Of course, further improvements are possible. For example the body part segmentations could be improved by a further application of Grabcut initialized from the current segmentations. Moreover, the model could be extended to include layering [1, 10, 24] to account for self-occlusions.

Acknowledgements. We thank Pawan Kumar for helpful discussions on probabilistic models, Varun Gulshan for the extended Grabcut, and Deva Ramanan for the image parsing code [18]. We are grateful for support from EU Project CLASS, the Royal Academy of Engineering, Microsoft, and the Spanish Ministry of Education and Science (grant FPU AP2003-2405, project TIN2005-01665).

References

- [1] A. Agarwal and B. Triggs. Tracking articulated motion using a mixture of autoregressive models. In *ECCV04*.
- [2] A. Agarwal and B. Triggs. 3d human pose from silhouettes by relevance vector regression. In *CVPR*, 2004.
- [3] M. Blank, L. Gorelick, E. Shechtman, M. Irani, and R. Basri. Actions as space-time shapes. In *ICCV*, 2005.
- [4] M. Bray, P. Kohli, and P. Torr. Posecut: Simultaneous segmentation and 3d pose estimation of humans using dynamic graph-cuts. In *ECCV*, 2006.
- [5] N. Dalal and B. Triggs. Histogram of oriented gradients for human detection. In *CVPR*, 2005.
- [6] P. Felzenszwalb and D. Huttenlocher. Pictorial structures for object recognition. *IJCV*, 61(1), 2005.
- [7] V. Ferrari, T. Tuytelaars, and L. Van Gool. Real-time affine region tracking and coplanar grouping. In *CVPR*, 2001.
- [8] N. Ikizler and P. Duygulu. Human action recognition using distribution of oriented rectangular patches. In *ICCV workshop on Human Motion Understanding*, 2007.
- [9] N. Jojic, J. Winn, and L. Zitnick. Escaping local minima through hierarchical model selection: Automatic object discovery, segmentation, and tracking in video. In *CVPR*, 2006.
- [10] M. P. Kumar, P. H. S. Torr, and A. Zisserman. Learning layered pictorial structures from video. In *ICVGIP*, pages 148–153, 2004.
- [11] M. P. Kumar, P. H. S. Torr, and A. Zisserman. Learning layered motion segmentations of video. In *ICCV*, 2005.
- [12] I. Laptev. Improvements of object detection using boosted histograms. In *BMV*, 2006.
- [13] I. Laptev and P. Perez. Retrieving actions in movies. In *ICCV*, 2007.
- [14] Z. Lin, L. Davis, D. Doermann, and D. DeMenthon. An interactive approach to pose-assisted and appearance-based segmentation of humans. In *ICCV workshop on Interactive Computer Vision*, 2007.
- [15] G. Mori, X. Ren, A. Efros, and J. Malik. Recovering human body configurations: Combining segmentation and recognition. In *CVPR*, 2004.
- [16] J. Niebles and L. Fei-Fei. A hierarchical model model of shape and appearance for human action classification. In *CVPR*, 2007.
- [17] M. Ozuysal, V. Lepetit, F. Fleuret, and P. Fua. Feature harvesting for tracking-by-detection. In *ECCV*, 2006.
- [18] D. Ramanan. Learning to parse images of articulated bodies. In *NIPS*, 2006.
- [19] D. Ramanan, D. A. Forsyth, and A. Zisserman. Strike a pose: Tracking people by finding stylized poses. In *CVPR*, 2005.
- [20] C. Rother, V. Kolmogorov, and A. Blake. Grabcut: Interactive foreground extraction using iterated graph cuts. In *SIGGRAPH*, 2004.
- [21] F. Schroff, A. Criminisi, and A. Zisserman. Single-histogram class models for image segmentation. In *ICVGIP*, 2006.
- [22] E. Shechtman and M. Irani. Matching local self-similarities across images and videos. In *CVPR*, 2007.
- [23] L. Sigal, S. Bhatia, S. Roth., M. Black, and M. Isard. Tracking loose-limbed people. In *CVPR*, 2004.
- [24] L. Sigal and M. Black. Measure locally, reason globally: Occlusion-sensitive articulated pose estimation. In *CVPR*, 2006.
- [25] J. Sivic, M. Everingham, and A. Zisserman. Person spotting: video shot retrieval for face sets. In *CIVR*, 2005.
- [26] C. Sminchisescu and B. Triggs. Estimating articulated human motion with covariance scaled sampling. *IJRR*, 2003.