

Detecting and Parsing of Visual Objects: Humans and Animals

Alan Yuille (UCLA)

Summary

- This talk describes recent work on detection and parsing visual objects. The methods represent objects in terms of object parts encoding spatial relations between them.
- We use deep convolutional neural networks (DCNNs) to make proposals for detecting the object parts.
- We use graphical models to encode spatial relationships between object parts and AND/OR graphs to share object parts between different, but similar, objects (e.g., cow torso and horse torso).
- We also describe AND/OR graphs for parsing humans.

Compositional Strategy

- Deep Convolutional Neural Networks (DCNNs) have been extremely successful for many visual tasks – such as object detection.
- But DCNNs are complicated “black boxes” and it is hard to understand what they are doing. They do not have explicit representations of object parts and the spatial relationships between them.
- Our strategy is to represent objects in terms of compositions of object parts. DCNNs are trained to detect parts. Then we use explicit graphical models – including AND/OR graphs – to encode spatial relations and to enable part sharing.
-

Organization – Three Parts

- Part (I). Parsing Humans – joint detection.
- Xianjie Chen and Alan Yuille (NIPS 2014, CVPR 2015).

- Part (II). Parsing Animals -- Semantic Segmentation.
- Peng Wang, Xiaohui Shen, Zhe Lin, Scott Cohen, Brian Price, Alan Yuille (ICCV 2015). Note: Shen, Lin, Cohen, Price are at Adobe Research.

- Part (III). Parsing Humans – Semantic Segmentation.
- Fangting Xia, Jun Zhu, Peng Wang, Alan Yuille (arxiv).


Part I: Parsing Human – Joint Detection



- In this project, the parts are joints (e.g., elbows, wrists, shoulders,...).
- Graphical models are used to represent spatial relationships between the parts.
- Part sharing is used to enable efficient inference when the human is occluded.
- X. Chen and A.L. Yuille (NIPS 2014, CVPR 2015).


Introduction

- ❑ Task is to estimate articulated human pose from a single static image.




□ Image Dependent Pairwise Relations (IDPRs)

- **Intuition:** We can reliably predict the relative positions of a part's neighbors (as well as the presence of the part itself) by only observing the local image patch around it.
- We specify a graphical model for human pose with novel pairwise relations that make adaptive use of local image measurements.



□ DCNN for Image Dependent Terms

- Require a method to extract information about pairwise part relations, as well as part presence, from local image patches.
- Deep Convolutional Neural Network (DCNN) is suitable for this, since it is efficient and share features between different parts and part relationships.



Performance Summary

□ State of the Art Performance

- Our model combines the representational flexibility of graphical models with the efficiency and statistical power of DCNNs.
- Significantly outperforms the state of the art methods on the **LSP** and **FLIC** datasets and also performs very well on the **Buffy** dataset without any training on it.

The Graphical Model

□ Variables of the Tree Model $\mathcal{G} = (\mathcal{V}, \mathcal{E})$

- The pixel location $\mathbf{l}_i = (x, y)$ of part $i \in \mathcal{V}$
- Pairwise relation types $t_{ij} \in \{1, \dots, T_{ij}\}, \forall (i, j) \in \mathcal{E}$

□ Unary Terms:

$$U(\mathbf{l}_i | \mathbf{I}) = w_i \phi(i | \mathbf{I}(\mathbf{l}_i); \boldsymbol{\theta})$$

□ Image Dependent Pairwise Relational (IDPR) Terms:


$$\begin{aligned} R(\mathbf{l}_i, \mathbf{l}_j, t_{ij}, t_{ji} | \mathbf{I}) = & \langle \mathbf{w}_{ij}^{t_{ij}}, \boldsymbol{\psi}(\mathbf{l}_j - \mathbf{l}_i - \mathbf{r}_{ij}^{t_{ij}}) \rangle + w_{ij} \varphi(t_{ij} | \mathbf{I}(\mathbf{l}_i); \boldsymbol{\theta}) \\ & + \langle \mathbf{w}_{ji}^{t_{ji}}, \boldsymbol{\psi}(\mathbf{l}_i - \mathbf{l}_j - \mathbf{r}_{ji}^{t_{ji}}) \rangle + w_{ji} \varphi(t_{ji} | \mathbf{I}(\mathbf{l}_j); \boldsymbol{\theta}) \end{aligned}$$

$$\boldsymbol{\psi}(\Delta \mathbf{l} = [\Delta x, \Delta y]) = [\Delta x \quad \Delta x^2 \quad \Delta y \quad \Delta y^2]^\top$$

□ The Full Score $F(\mathbf{l}, \mathbf{t} | \mathbf{I}) = \sum_{i \in \mathcal{V}} U(\mathbf{l}_i | \mathbf{I}) + \sum_{(i,j) \in \mathcal{E}} R(\mathbf{l}_i, \mathbf{l}_j, t_{ij}, t_{ji} | \mathbf{I}) + w_0$

DCNN for Image Dependent Terms

- ▣ Appearance terms $\phi(\cdot|\cdot; \theta)$ and IDPR terms $\varphi(\cdot|\cdot; \theta)$ depend on the image.
- ▣ We use DCNN to learn the conditional probability distribution $p(c, \mathbf{m}_{c\mathcal{N}(c)} | \mathbf{I}(\mathbf{l}_i); \theta)$ defined on the space $|\mathcal{S}|$, where each element corresponds to a part with all the types of its pairwise relationships, or the background.



- ▣ Marginalization: $\phi(i | \mathbf{I}(\mathbf{l}_i); \theta) = \log(p(c = i | \mathbf{I}(\mathbf{l}_i); \theta))$ $\varphi(t_{ij} | \mathbf{I}(\mathbf{l}_i); \theta) = \log(p(m_{ij} = t_{ij} | c = i, \mathbf{I}(\mathbf{l}_i); \theta))$

Inference

- ❑ Dynamic programming + Distance Transform: $O(T^2 LK)$
 - L: # of locations, K: # of parts, T: # of pairwise types
- ❑ Image Dependent Terms are efficiently calculated by a single DCNN at all locations.
 - The computations common to overlapping regions are shared by considering fully-connected layers as 1x1 convolutions.

Relationship to other models

❑ Pictorial Structure (PS)

- Recover PS by allowing one pairwise relation type, i.e., $T_{ij} = 1$
- We use DCNN to learn data term instead of HOG filters.

❑ Yang and Ramanan's Mixtures-of-parts (MOP) [26]

- MOP defines different “types” of part by its relative position with respect to its parent.
- Recover MOP by restricting each part in our model to only predict the relative position of its parent, i.e., $T_{ij} = 1$ if j is not parent of i .

❑ Conditional Random Fields (CRFs)

- Related to CRFs literature on data dependent priors.
- Efficiently model all the image dependent terms in a single DCNN.

Learning

- ❑ Supervised learning by deriving pairwise type labels from the annotated part (joint) locations by clustering.
- ❑ Learn three sets of parameters:
 - Mean relative positions \mathbf{r} of different pairwise relation types, by K-means clustering.
 - Parameters θ of image dependent terms, by DCNN.
 - Weight parameters \mathbf{W} , by linear SVM.

Benchmark Performance

□ LSP

Method	Torso	Head	U.arms	L.arms	U.legs	L.legs	Mean
Ours	92.7	87.8	69.2	55.4	82.9	77.0	75.0
Pishchulin et al. [16]	88.7	85.6	61.5	44.9	78.8	73.4	69.2
Ouyang et al. [14]	85.8	83.1	63.3	46.6	76.5	72.2	68.6
DeepPose* [23]	-	-	56	38	77	71	-
Pishchulin et al. [15]	87.5	78.1	54.2	33.9	75.7	68.0	62.9
Eichner&Ferrari [4]	86.2	80.1	56.5	37.4	74.3	69.3	64.3
Yang&Ramanan [26]	84.1	77.1	52.5	35.9	69.5	65.6	60.8

Table 1: Comparison of *strict* PCP results on the LSP dataset. Our method improves on all parts by a significant margin, and outperforms the best previously published result [1] by 5.8% on average. Note that DeepPose uses Person-Centric annotations and is trained with an extra 10,000 images.

- Two recent ECCV'14 papers, Kiefel&Gehler and Ramakrishna et al., also report performance on the LSP dataset, and our performance is better.

FLIC

Method	U.arms	L.arms	Mean
Ours	97.0	86.8	91.9
MODEC [20]	84.4	52.1	68.3

Table 2: Comparison of *strict* PCP results on the FLIC dataset. Our method significantly outperforms MODEC [20].





Figure 1: Comparison of PDJ curves of elbows and wrists on the FLIC dataset. The legend shows the PDJ numbers at the threshold of 0.2.

Datasets

- ❑ **Leeds Sports Poses (LSP) dataset:** 1000 training and 1000 testing full-body human poses.
- ❑ **Frames Labeled In Cinema (FLIC) dataset:** 3987 training and 1016 testing upper-body human poses.
- ❑ **Buffy Stickmen dataset:** 276 testing upper-body human poses. *We do not train on this dataset.*



Diagnostic Experiments

□ Terms Analysis

Method	Torso	Head	U.arms	L.arms	U.legs	L.legs	Mean
<i>Unary-Only</i>	56.3	66.4	28.9	15.5	50.8	45.9	40.5
<i>No-IDPRs</i>	87.4	74.8	60.7	43.0	73.2	65.1	64.6
Full Model	92.7	87.8	69.2	55.4	82.9	77.0	75.0

Table 3: Diagnostic term analysis *strict* PCP results on the LSP dataset. The unary term alone is still not powerful enough to get good results, even though it's trained by a DCNN classifier. *No-IDPRs* method, whose pairwise terms are not dependent on the image, can get comparable performance with the state-of-the-art, and adding IDPR terms significantly boost our final performance to 75.0%.

□ Cross-dataset Generalization

Method	U.arms	L.arms	Mean
Ours*	96.8	89.0	92.9
Ours* <i>strict</i>	94.5	84.1	89.3
Yang [27]	97.8	68.6	83.2
Yang [27] <i>strict</i>	94.3	57.5	75.9
Sapp [21]	95.3	63.0	79.2
FLPM [11]	93.2	60.6	76.9
Eichner [5]	93.2	60.3	76.8

Table 3: Cross-dataset PCP results on Buffy test subset. The PCP numbers are *Buffy* PCP unless otherwise stated.




Figure 2: Cross-dataset PDJ curves on Buffy test subset. The legend shows the PDJ numbers at the threshold of 0.2. Note that both our method and DeepPose [23] are trained on the FLIC dataset.

- Compared with Figure 1., the margin between our method and DeepPose significantly increases in Figure 2., which implies that our model generalizes better to the Buffy dataset.

Parsing People by Flexible Compositions. (Chen and Yuille CVPR 2015).

- In realistic images many object parts are occluded.
- Previous graphical model are robust to only a few occlusions.
- Prior – observed nodes of graphical model are often connected.
- Strategy: extend the method used in NIPS 2014 to deal with occlusion.





Figure 1: An illustration of the *flexible compositions*. Each connected subtree of the *full graph* (include the full graph itself) is a flexible composition. The flexible compositions that do not have certain parts is suitable for the people with those parts occluded.




(a)



(b)

Figure 2: Motivation. (a): In real world scenes, people are usually significantly occluded (or truncated). Requiring the model to localize a fixed set of body parts while ignoring the fact that different people have different degrees of occlusion (or truncation) is problematic. (b): The absence of body parts evidence can help to predict occlusion, *e.g.*, the right wrist of the lady in brown can be inferred as occluded because of the absence of suitable wrist near the elbow. However, absence of evidence is not evidence of absence. It can fail in some challenging scenes, for example, even though the left arm of the lady in brown is completely occluded, there is still strong image evidence of suitable elbow and wrist at the plausible locations due to the confusion caused by nearby people (*e.g.*, the lady in green). In both situations, the local image measurements near the occlusion boundary (*i.e.*, around the right elbow and left shoulder), *e.g.*, in a image patch, can reliably provide evidence of occlusion.



- Elbow: ●
- Lower Arm: 
- Upper Arm: 
- Occluders: 

Figure 3: Different occlusion decoupling and spatial relationships between the elbow and its neighbors, *i.e.*, wrist and shoulder. The local image measurement around a part (*e.g.*, the elbow) can reliably predict the relative positions of its neighbors when they are not occluded, which is demonstrated in the base model [5]. In the case when the neighboring parts are occluded, the local image measurement can also reliably provide evidence for the occlusion.

Model

- Base Model: as before.
- Introduce decoupling terms
- Penalties for missing terms

$$D_{ij}(\gamma_{ij} = 1, \mathbf{l}_i | \mathbf{I}) = w_{ij} \varphi^d(\gamma_{ij} = 1 | \mathbf{I}(\mathbf{l}_i); \boldsymbol{\theta}),$$

$$B_{ij} = \sum_{k \in \mathcal{V}(T_j)} b_k$$

$$\begin{aligned} F(\mathbf{l}, \mathbf{t}, \mathcal{G}_c | \mathbf{I}, \mathcal{G}) &= \sum_{i \in \mathcal{V}_c} A(\mathbf{l}_i | \mathbf{I}) \\ &+ \sum_{(i,j) \in \mathcal{E}_c} R(\mathbf{l}_i, \mathbf{l}_j, t_{ij}, t_{ji} | \mathbf{I}) \\ &+ \sum_{(i,j) \in \mathcal{E}_c^d} (B_{ij} + D_{ij}(\gamma_{ij} = 1, \mathbf{l}_i | \mathbf{I})) \end{aligned} \quad (5)$$

Inference

- There are many different models – no. of connected subtrees of the graph.
- But inference is efficient because of part-sharing.
- Inference is only twice the complexity of the base model:
- $O(2T^2LK)$

Evaluation

- “We Are Family” (WAF) Dataset
- 525 images, six people per image on average. (350/175 train/test).

Method	AOP	Torso	Head	U.arms	L.arms	mPCP
Ours	84.9	88.5	98.5	77.2	71.3	80.7
Multi-Person [9]	80.0	86.1	97.6	68.2	48.1	69.4
Ghiasi et. al. [15]	74.0	-	-	-	-	63.6
One-Person [9]	73.9	83.2	97.6	56.7	28.6	58.6

Table 1: Comparison of PCP and AOP on the WAF dataset. Our method improves the PCP performance on all parts, and significantly outperform the best previously published result [9] by 11.3% on mean PCP, and 4.9% on AOP.

Diagnostics

Method	AOP	Torso	Head	U.arms	L.arms	mPCP
Base Model [5]	73.9	81.4	92.6	63.6	47.6	66.1
<i>FC</i>	82.0	87.0	98.6	72.7	67.5	77.7
<i>FC+IDOD</i>	84.9	88.5	98.5	77.2	71.3	80.7

Table 2: Diagnostic Experiments PCP and AOP results on the WAF dataset. Using flexible compositions (*i.e.*, *FC*) significantly improves our base model [5] by 11.6% on PCP and 8.1% on AOP. Adding *IDOD* terms (*FC+IDODs*, *i.e.*, the full model) further improves our PCP performance to 80.7% and AOP performance to 84.9%, which is significantly higher than the state of the art methods.




Figure 5: Results on the WAF dataset. We show the parts that are inferred as visible, and thus have estimated configurations, by our model.

From 2D to 3D.

- Pose detection – with and with occlusion.
 - Prior – connected parts – for occlusion (validated on WAF)
 - Efficient inference despite occlusion – due to part sharing.
- Note: detection of pose is important for many applications.
E.g., estimating of 3D structure (C. Wang et al. 2014), action recognition (C. Wang et al, 2013, 2014).

Collaboration with Peking University.



Summary of Part I: Parsing Humans -- Joints

- Detection of object parts (joints) in presence of occlusion. DCNNs for detecting parts, graphical models to impose spatial relations, efficient inference using dynamic programming.
- The detected parts can be used to estimate 3D structure of humans from a single image and enable action recognition.
- Limitations. Objects are represented in terms of joints only. This becomes problematic in some human configurations.

Part II – Parsing Animals – Semantic Segmentation

- Detecting and Parsing of Animals. Semantic Segmentation.
- The parts are heads, arms, torsos, legs, tails.
- The parts are shared between different animals.
- We perform semantic segmentation – i.e. labeling the pixels of each part.
- P. Wang, X. Shen, Z. Lin, S. Cohen, B. Price, A. Yuille. ICCV 2015.

The Task


Original
image



Object mask



Part mask




Detect and Parse Animals:
Jointly infer the object segmentation and part
segments

Motivation


- Detecting parts can improve object detection:
 - Strongly supervised DPM [Azizpour et.al ECCV 2012]
 - Detect what you can [X. Chen et.al CVPR 2014]
 - ..
- The same intuitions apply to segmentation,
 - Parts need less context and can provide finer boundaries
 - Object needs long range context, but miss details.
 - *Parsing and segmenting objects in term of parts give rich descriptions suitable for many visual tasks.*

The Framework



- Our method is performed using two stages:
- Object and Part potentials -- make proposals for parts.
- Fully connected CRF – spatial relations and part sharing.

Part Sharing: Semantic compositional parts (SCP)




We use part sharing to reduce the amount of computation.

$l_{op} \in \{\text{horse-head, horse-body, horse-leg, horse-tail, cow-head, cow-leg, cow-body, cow-tail}\}$

SCP segments proposal

□ Parts are detected despite the difficulty of the data.



Joint FCRF: representing spatial relations


$$\min_{\mathcal{L}} \sum_{i \in \mathcal{V}} \psi_i(l_{op}^i) + \lambda_e \sum_{i, j \in \mathcal{V}, i \neq j} \psi_{i,j}(l_{op}^i, l_{op}^j)$$

where, $\psi_i(l_{op}^i) = \eta(l_o^i, l_p^i)(\psi_i^o(l_o^i)) + \lambda_p \psi_i^p(l_p^i); \quad (1)$

$$\psi_{i,j}(l_{op}^i, l_{op}^j) = \eta(l_o^i, l_p^i) \eta(l_o^j, l_p^j) \psi_{i,j}^{op}(l_o^i, l_o^j, l_p^i, l_p^j);$$

$$\psi_i^o(l_o^i) = \sum_{\mathbf{x}_j \in \mathbf{P}_i} -\log(P(l_o(\mathbf{x}_j)))$$

$$\psi_{i,j}^{op}(l_o^i, l_o^j, l_p^i, l_p^j) = -\log(P(l_o^i)P(l_o^j)P(l_p^i)P(l_p^j))$$



Two layer neural network, with features from

- Context potential
- Relative spatial position
- Geodesic and Euclidean distance

Experiments

- Data (3 Dataset derived from our PASCAL part [X. Chen et. al CVPR 2014])
 - Horse-Cow Data
 - The Quadrupeds Data
 - Pascal part benchmark
 -
- Comparison
 - Semantic part segmentation (SPS) [J. Wang and A.L. Yuille CVPR 2015]
 - Hypercolumn [Hariharan et.al CVPR 2015]
 - FCN for object segmentation [Long et. al CVPR 2015]

□ Horse-Cow Data

□ Data from SPS, segment horse and cow parts given bounding box.

Horse								
	Bkg	head	body	leg	tail	Fg	IOU	Pix. Acc
SPS [39]	79.14	47.64	69.74	38.85	-	68.63	-	81.45
HC [21]	85.71	57.30	77.88	51.93	37.10	78.84	61.98	87.18
Ours	87.34	60.02	77.52	58.35	51.88	80.70	65.02	88.49

Cow								
	Bkg	head	body	leg	tail	Fg	IOU	Pix. Acc
SPS [39]	78.00	40.55	61.65	36.32	-	71.98	-	78.97
HC [21]	81.86	55.18	72.75	42.03	11.04	77.07	52.57	84.43
Ours	85.68	58.04	76.04	51.12	15.00	82.63	57.18	87.00

Table 1. Average precision over the Horse-Cow dataset.

□ The Quadrupeds Data

□ 5 animal classes from PASCAL

Object segmentation accuracy								
	Bkg	Dog	Cat	Cow	Horse	Sheep	IOU	Pix. Acc
FCN 16s [27]	93.25	74.30	78.62	61.88	56.56	67.63	72.04	93.00
FCN 8s [27]	93.55	74.39	78.52	60.81	58.39	69.15	72.47	93.17
Joint FCN(16s)	94.04	75.13	80.52	66.76	63.04	71.54	75.17	93.77
FCRF+FCN(16s)	93.88	77.10	80.92	68.76	63.40	64.54	74.57	93.87
Ours final	94.40	79.03	83.04	74.82	69.94	70.59	78.64	94.71

Semantic part segmentation accuracy								
	Bkg	Dog	Cat	Cow	Horse	Sheep	IOU	Pix. Acc
HC [21]	92.83	42.07	43.99	35.49	38.59	33.80	41.36	89.54
Ours final	94.46	45.63	47.81	42.7	49.60	35.74	46.69	91.74

Table 2. Average precision over the Quadrupeds data.

□ Pascal part benchmark


□ Quadrupeds part segments from Pascal test set

Object segmentation accuracy								
	Bkg	Dog	Cat	Cow	Horse	Sheep	IOU	Pix. Acc
FCN 8s [27]	94.45	70.14	75.45	64.06	64.75	69.06	72.99	93.90
Ours final	95.31	77.44	80.47	72.13	76.18	67.96	78.25	95.26

Semantic part segmentation accuracy								
	Bkg	Dog	Cat	Cow	Horse	Sheep	IOU	Pix. Acc
HC [21]	94.36	41.24	42.42	35.22	45.00	38.86	43.11	90.64
Ours final	95.14	46.52	48.06	41.80	56.67	36.02	48.16	92.47

Table 3. Average precision over the part segmentation benchmark.

Qualitative results



Summary Part 2: Animal Parsing, Semantic Segmentation

Detect and semantically segment object parts.

Limitations: Occlusions, Small Ambiguous Parts (e.g., Tails).


Part III: Parsing Humans – Semantic Segmentation. Pose-Guided Human Body Parsing with Deep-Learned Features

Fangting Xia, Jun Zhu, Peng Wang and Alan Yuille


Motivation




(a)



(b)




(c)




(d)


The Human Parsing Pipeline




The AOG model for part assembling



The AOG Model



(a) leaf vertex



(b) non-leaf vertex

$$F(\mathbf{Y}, \mathbf{Z} | \tilde{\mathcal{S}}, \mathcal{L}, \mathcal{H}) = \sum_{p \in \mathcal{T}} f(y_p, z_p) \quad (4)$$

$$+ \sum_{c \in \mathcal{N}} f(y_c, z_c, \{(y_\mu, z_\mu) : \mu \in Ch(c, z_c)\}),$$

$$f(y_p, z_p) = \begin{cases} b_{z_p}^p + w_{z_p}^p \cdot g_{z_p}^p(s_{y_p}^{p, z_p} | \mathcal{L}, \mathcal{H}), & z_p \neq 0 \\ b_0^p, & z_p = 0 \end{cases} \quad (5)$$

$$f(y_c, z_c, \{(y_\mu, z_\mu) : \mu \in Ch(c, z_c)\}) \quad (6)$$

$$= \begin{cases} b_{z_c}^c + u(y_c, z_c, \{(y_\mu, z_\mu) : \mu \in Ch(c, z_c)\}), & z_c \neq 0 \\ b_0^c, & z_c = 0 \end{cases}$$

where

$$u(y_c, z_c, \{(y_\mu, z_\mu) : \mu \in Ch(c, z_c)\}) \quad (7)$$

$$= h(\{(y_\mu, z_\mu) : \mu \in Ch(c, z_c)\})$$


$$+ \sum_{\mu \in Ch(c, z_c)} \mathbf{w}_{(z_c, z_\mu)}^{(c, \mu) T} \varphi(s_{y_c}^{c, z_c}, s_{y_\mu}^{\mu, z_\mu})$$

$$+ \sum_{(p_1, p_2) \in \mathcal{R}_c} \mathbf{w}_{(z_{p_1}, z_{p_2})}^{(p_1, p_2) T} \psi(s_{y_{p_1}}^{p_1, z_{p_1}}, s_{y_{p_2}}^{p_2, z_{p_2}} | \mathcal{L}),$$


$$h(\{(y_\mu, z_\mu) : \mu \in Ch(c, z_c)\}) = \sum_{\mu \in Ch(c, z_c) \cap \mathcal{T}} f(y_\mu, z_\mu) \quad (8)$$

$$+ \sum_{\mu \in Ch(c, z_c) \cap \mathcal{N}} f(y_\mu, z_\mu, \{(y_\nu, z_\nu) : \nu \in Ch(\mu, z_\mu)\}).$$


The Pose Context Feature



(a)



(b)



(c)


The Unary Part Prototypes Learned from Pose Context Feature



The Pairwise Part Prototypes Learned From Pose Context Feature



The Effect of Pose Cues for Part Segment Proposal Generation



The Recall and Average IoU of Our Segment Proposals

	hair	face	u-cloth	l-cloth	arms	legs	shoes	mean
Recall	0.88	0.90	0.99	0.99	0.86	0.87	0.67	0.88
IoU	0.73	0.74	0.85	0.86	0.67	0.72	0.56	0.73

The Effect of Various Features (Comb.) for Part Proposal Ranking and Selection

	hair	face	u-cloth	l-cloth	arms	legs	shoes	mean
o2p+skin	57.1	53.5	70.9	70.9	26.6	20.4	15.6	45.0
	68.8	66.9	80.0	81.4	54.6	55.3	45.3	64.6
(1)+PCF	61.7	58.6	73.2	72.7	29.9	23.4	17.5	48.1
	69.9	66.4	80.6	82.3	56.4	54.3	45.8	65.1
(2)+C- PCF	61.8	58.9	73.2	71.9	39.8	44.8	26.5	53.8
	69.9	66.4	80.5	82.4	55.8	59.1	47.4	65.9
(3)+deep potential	64.4	59.0	77.4	77.1	41.4	43.6	35.1	56.9
	70.7	66.6	82.2	83.4	55.9	59.3	48.8	66.7

Top-1 (upper row) and top-10 (lower row) AOI scores of part ranking models

Investigating The Effect of AOG for Part Assembling

	hair	face	u-cloth	arms	l-cloth	legs	Avg
Naive assembling	62.3	53.5	77.8	36.9	78.3	28.2	56.2
Basic AOG	63.1	52.9	77.1	38.0	78.1	35.9	57.5
Ours	63.2	56.2	78.1	40.1	80.0	45.5	60.5
Ours (w/o pruning)	63.2	56.2	78.1	40.1	80.0	45.8	60.5

Naive assembling: using only the unary terms and basic geometric common sense constraints (e.g. upper-clothes and lower-clothes must be adjacent).

Basic AOG: using only the unary terms and the parent-child pairwise terms, without the side-way pairwise terms between parts.

Ours: using all potentials together (including the unary terms, the parent-child pairwise terms, and the side-way pairwise terms).

Ours (w/o pruning): the results of our model without greedy pruning. The pruning only brings ignorable decrease in performance while it reduces the inference time from 2 min. to 1 sec. per image.

Comparison to The State of The Art

	hair	face	u-cloth	arms	l-cloth	legs	shoes	Avg*
FCN-32 [17]	50.2	33.7	69.4	13.8	66.7	14.2	25.2	41.3
FCN-16 [17]	48.7	49.1	70.2	33.9	69.6	29.9	36.1	50.2
P&S [20]	40.0	42.8	75.2	24.7	73.0	46.6	-	50.4
SBP [2]	44.9	60.8	74.8	26.2	71.2	42.0	-	53.3
DDN [18]	43.2	57.1	77.5	27.4	75.3	52.3	-	56.2
Ours	63.2	56.2	78.1	40.1	80.0	45.5	35.0	60.5

Comparison of our approach with other state-of-the-art algorithms over the Penn-Fudan dataset. The Avg means the average without shoes class since it did not included in other algorithms.


Qualitative Results



Qualitative Result Comparison Between Our Method and FCNN



Image



FCN-32



FCN-16



Ours

Some Failure Cases of Our Method



Failure cases of our algorithm on Penn-Fudan dataset. For each case, the original image (with pose prediction), ground truth, and our parsing result are displayed from left to right.

Summary of Part III: Human parsing – semantic segmentation

- Parsing humans is more difficult than parsing animals – because humans wear clothes, and there are many choices of clothes.
- Our approach uses deep networks for joints (X. Chen and A.L. Yuille), pose-context features, appearance cues (including deep networks for large parts).
- AND/OR graph is used to allow us to model the large number of possible configurations of humans.

Summary

- This talk illustrated a research program where we represent objects as compositions of object parts.
- We use deep convolutional neural networks (DCNNs) to make proposals for the parts. The parts can either be human joint (e.g., elbows) or animal parts (e.g., head and torso).
- Graphical model are used to capture spatial relations between object parts and to enable part sharing (e.g., horse torso and cow torso).
- This approach gives state of the art results on benchmarked datasets.

Papers Cited.

- X. Chen and A.L. Yuille. Articulated Pose Estimation with Image-Dependent Preference on Pairwise Relations. NIPS 2014.
- X. Chen and A.L. Yuille. Parsing Occluded People by Flexible Compositions. CVPR 2015.
- P. Wang, P. Wang, X. Shen, Z. Lin, S. Cohen, B. Price, A. Yuille. Joint Object and Part Segmentation using Deep Learned Potentials. ICCV 2015.
- Z. Xia, J. Zhu, P. Wang, A.L. Yuille. Pose-Guided Human Body Parsing with Deep-Learned Features. Arxiv. 2015.
- C. Wang, Y. Wang, Z. Lin, A.L. Yuille, and W. Gao. Robust Estimation of 3D Human Poses from Single Images. CVPR. 2014.