Correlated Uncertainty for Learning Dense Correspondences from Noisy Labels

Natalia Neverova, David Novotny, Andrea Vedaldi

Facebook AI Research

facebook Artificial Intelligence

Dense Pose

Label pixels with their location on a surface of a reference 3D human body model

Learning Dense Pose

Input image I

Deep net

Predicted UV

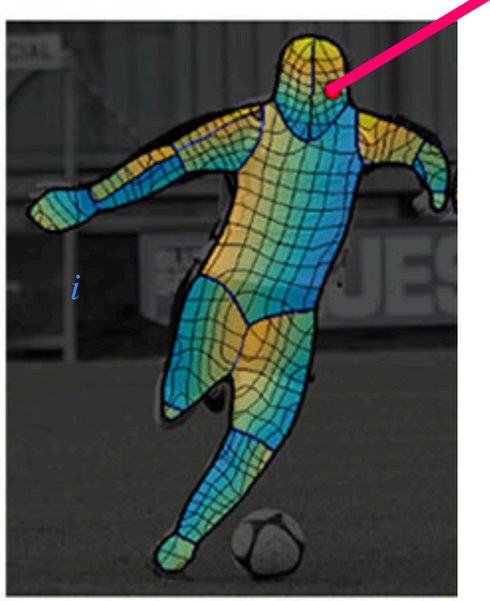

Pixel i

GT label $y_i \in R^2$

prediction $\hat{y}_i \in R^2$

residual $\delta_i = y_i - \hat{y}_i$

Naive approach:

$$\ell = \sum_i \|\delta_i\|^2$$

Dense Pose annotations are noisy

Annotation process:

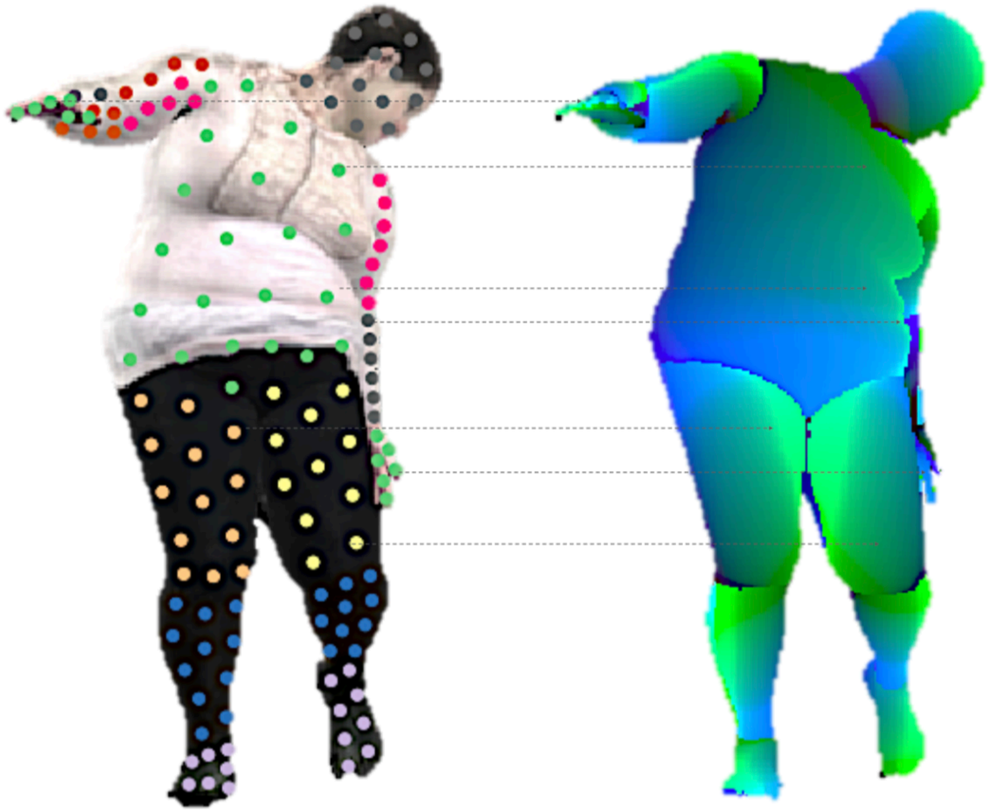

Manual annotations

Rendered image

Dense Pose annotations are noisy

Annotations projected
on the 3D model:

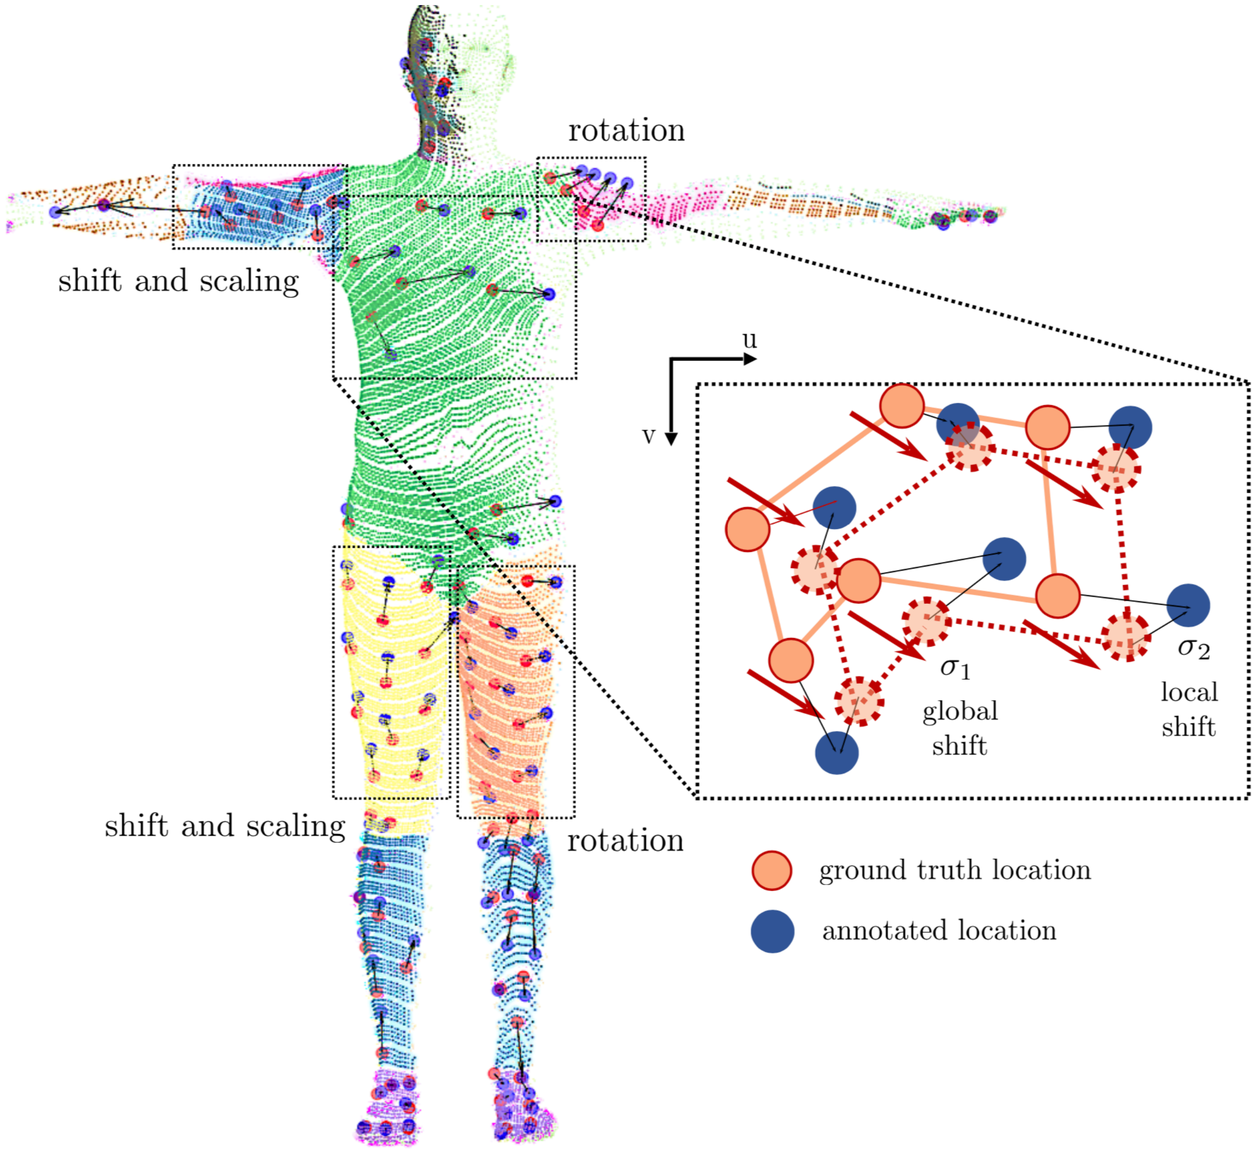

Elementary uncertainty model

Input image I

Deep net

Predicted UV

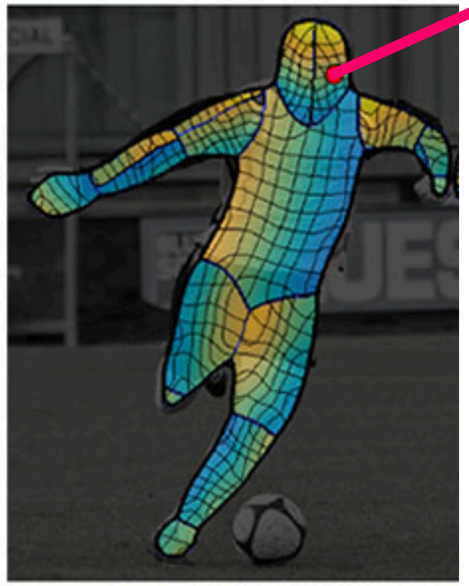

Model $p(\delta_i)$ for every pixel i [1,2]:

normal distr. of residuals δ_i:

$$p(\delta_i) = \mathcal{N}(\delta_i | 0, \sigma_i)$$

network predictions:

$\hat{y}_i \in R^2$ and $\sigma_i \in R$

maximize log-likelihood:

$$\ell = -\sum_i \log N(\delta_i | 0, \sigma_i)$$

1) Isotropic variance - non-directional errors
2) IID assumption on errors

[1] Novotny et al.: Learning 3D Object Categories By Looking Around Them
[2] Kendall & Gal: What Uncertainties Do We Need in Bayesian Deep Learning for Computer Vision?

Higher-order uncertainty model

Model $p(\vec{\delta})$ for the whole part:

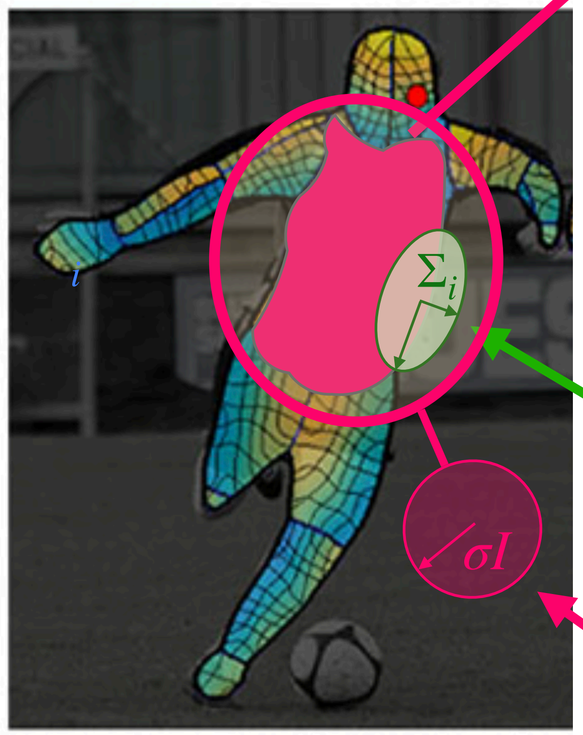

$p(\vec{\delta})$ gaussian: $\Sigma = \begin{bmatrix} \Sigma_1 + \sigma I & \sigma I & \dots & \sigma I \\ \sigma I & \Sigma_2 + \sigma I & \dots & \sigma I \\ \vdots & \ddots & \ddots & \vdots \\ \sigma I & \sigma I & \cdots & \Sigma_N + \sigma I \end{bmatrix}$

1) Directional distr. of residuals δ_i:

$$p(\delta_i) = \mathcal{N}(\delta_i | 0, \Sigma_i) \qquad \Sigma_i \in R^{2\times 2}$$

network prediction

2) Overall part-specific offset ϵ

$$p(\epsilon) = \mathcal{N}(0, \sigma I)$$

... models overall correlation between per-pixel error directions

Experiments

Performance on DensePose-COCO for different error thresholds

uv-loss	1 cm	2 cm	3 cm	5 cm	10 cm	20 cm
MSE	4.44	16.21	29.64	52.23	76.50	**85.99**
full	**5.99**	**19.97**	**34.16**	**55.68**	**77.76**	85.58

Ours

Experiments

Ablation of the probabilistic terms

LogLikelihood attained on test set of DensePose-COCO

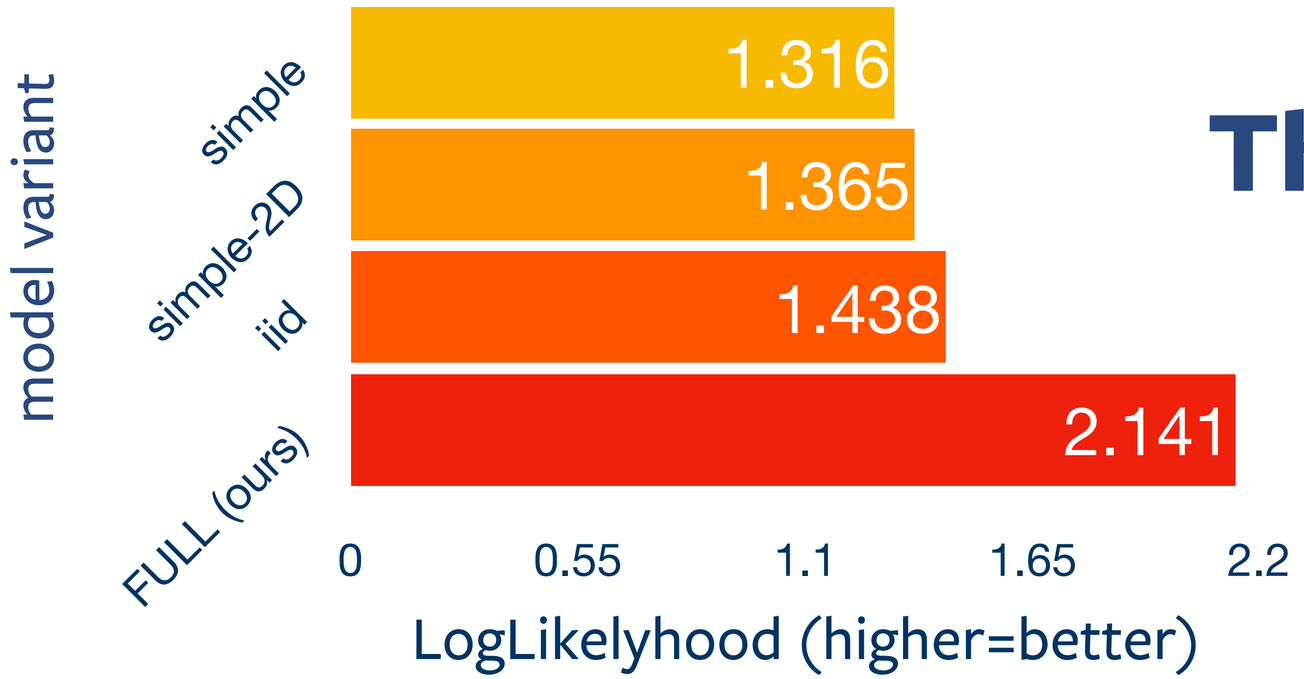

Thank you!