

Intuitions

- Pose normalization: predict keypoints, prioritize those locations
- Used to work well on fine-grained classification

(Farrell et al. 2011)

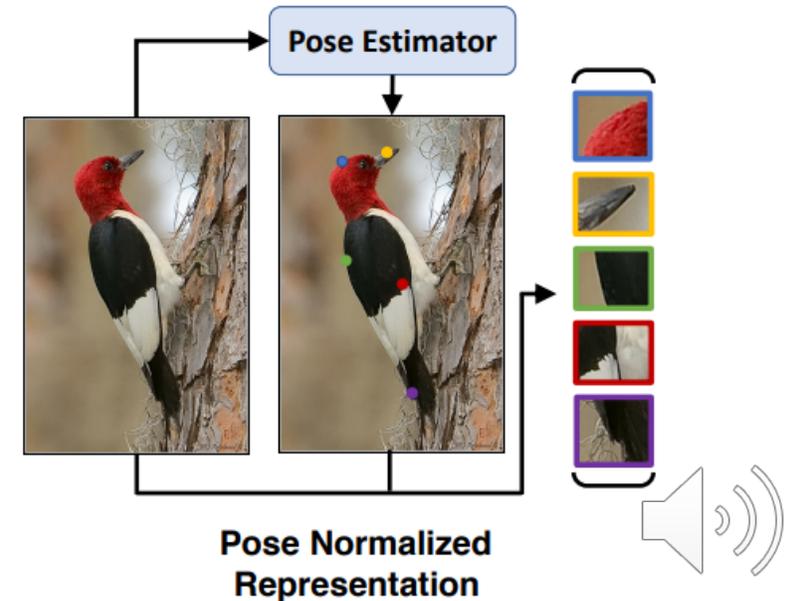
- But end-to-end black-box networks work better (e.g. bilinear pooling)
- But is this true for **few-shot** settings?

the Task:
red headed woodpecker



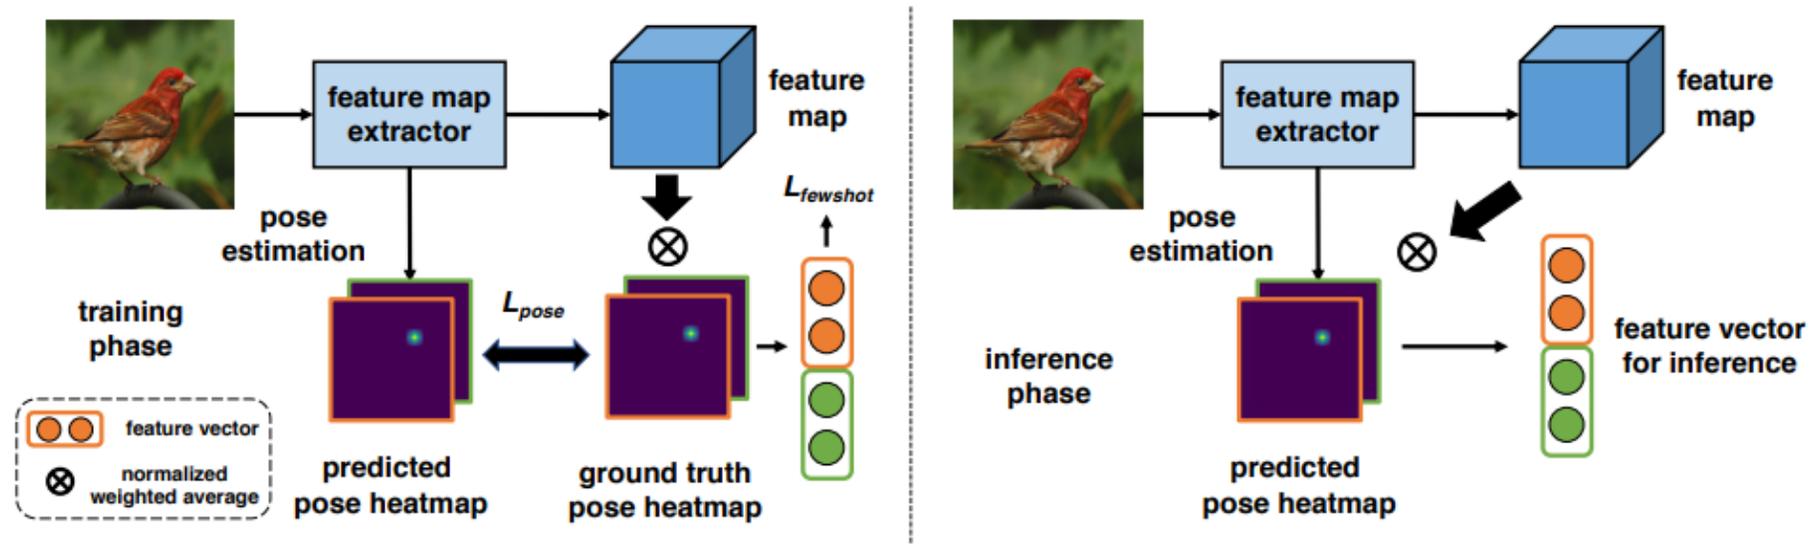
red bellied woodpecker

?



Pose Normalization Model

- 2-layer convolutional pose estimation module, with very few parameters added
- Require no part annotation for novel categories
- Easily adopted to few-shot learning algorithms, plug and play



Results

- Improves accuracy between 10 and 20 percentage points.
- Robust to part annotation sparsity.

