Local Propagation for Few-Shot Learning

Yann Lifchitz[1,2], Yannis Avrithis[1], Sylvaine Picard[2]

[1]Univ Rennes, Inria, CNRS, IRISA [2]Safran

Transductive Few-shot learning

- Classify queries into previously unseen classes with very few annotated examples
- Transduction: queries are seen in batch
- Transductive methods leverage their distribution for better classification

Contributions

- Study of graph-based propagation of on local representations across images
- Application to few-shot learning, bridging the gap between transductive and non-transductive inference
- Introduction of a simple but powerful spatial attention mechanism

Process

- Base class training using dense classification [2]
- Spatial attention and extraction of local features of supports and queries
- Building a graph with supports and queries local features
- Optionally propagating features on the graph
- Classification of queries with label propagation

Local Features

- Pixels of feature maps are treated as local features
- **Spatial attention**: removes local features corresponding to irrelevant part of the image
- $a(F) := \{F(r) : \|F(r)\| \geq \tau \max_{t \in \Omega} \|F(t)\|, r \in \Omega\}$, $F(r)$ being the feature vector corresponding to spatial position r in the position set Ω.
- **Feature pooling**: Clustering of the remaining features using k-means clustering
- Final local features are the resulting cluster centroids

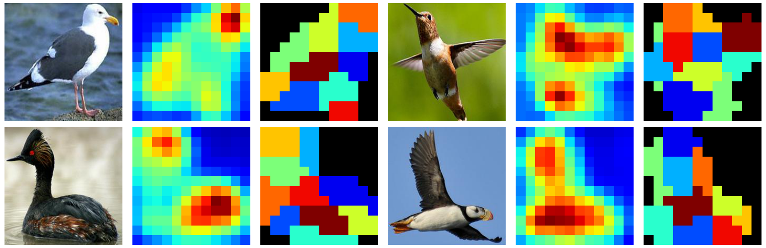

Examples of CUB images, each with the corresponding spatial attention heatmap and clusters used in feature pooling.

Local Propagation

- Vertices: Local features
- Edges values: Cosine similarity
- Propagation as in [7]
- Propagated labels treated as predicted class probability
- Predictions average for each query

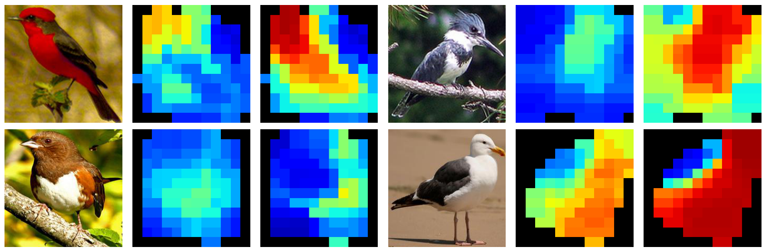

Examples of CUB query images in 5-way 5-shot non-transductive tasks, each followed by the heatmap of predicted probability for the correct class using a prototype classifier, then using local label propagation.

Results

Method	CUB 1-shot	CUB 5-shot	*mini*ImageNet 1-shot	*mini*ImageNet 5-shot
Non-transductive inference				
Proto [5]	74.85±0.48	90.38±0.27	63.39±0.46	81.21±0.32
Proto [5]+Att	77.10±0.47	91.24±0.26	64.22±0.45	81.71±0.31
Global label propagation	77.23±0.46	88.78±0.31	63.41±0.45	77.04±0.37
Local label propagation	79.32±0.44	91.52±0.25	64.43±0.45	80.26±0.32
Transductive inference				
TPN [3]	-	-	59.46	75.65
LR+ICI [6]	88.06	92.53	66.80	79.26
EPNet [4]	82.85±0.81	91.32±0.41	66.50±0.89	81.06±0.60
Global label propagation	87.18±0.46	91.88±0.27	72.54±0.54	81.38±0.35
Local label propagation	87.77±0.41	93.35±0.23	72.57±0.51	82.76±0.33

5-way few-shot classification accuracy. All propagation methods use spatial attention, feature propagation and feature pooling if possible.

Universal solution

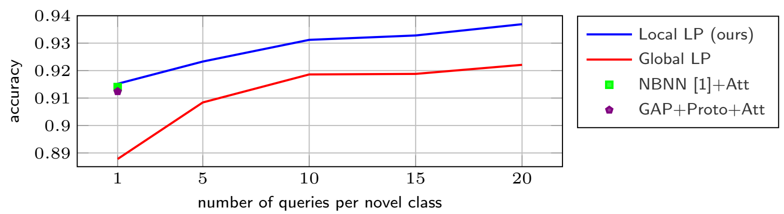

CUB 5-way 5-shot classification accuracy *vs.* number of queries per novel class.

References

[1] W. Li, L. Wang, J. Xu, J. Huo, Y. Gao, and J. Luo. Revisiting local descriptor based image-to-class measure for few-shot learning. In *CVPR*, 2019.

[2] Y. Lifchitz, Y. Avrithis, S. Picard, and A. Bursuc. Dense classification and implanting for few-shot learning. *CVPR*, 2019.

[3] Y. Liu, J. Lee, M. Park, S. Kim, E. Yang, S. Hwang, and Y. Yang. Learning to propagate labels: Transductive propagation network for few-shot learning. In *ICLR*, 2019.

[4] P. Rodríguez, I. Laradji, A. Drouin, and A. Lacoste. Embedding propagation: Smoother manifold for few-shot classification. *arXiv preprint arXiv:2003.04151*, 2020.

[5] J. Snell, K. Swersky, and R. Zemel. Prototypical networks for few-shot learning. In *NIPS*, 2017.

[6] Y. Wang, C. Xu, C. Liu, L. Zhang, and Y. Fu. Instance credibility inference for few-shot learning. *arXiv preprint arXiv:2003.11853*, 2020.

[7] D. Zhou, O. Bousquet, T. N. Lal, J. Weston, and B. Schölkopf. Learning with local and global consistency. In *NIPS*, 2003.

Contact: yann.lifchitz@safrangroup.com