LOCAL PROPAGATION FOR FEW-SHOT LEARNING

—

Yann Lifchitz[1,2], Yannis Avrithis[1], Sylvaine Picard[2]
[1]Inria, Univ Rennes, CNRS, IRISA [2]Safran

Inria INVENTEURS DU MONDE NUMÉRIQUE SAFRAN

Introduction of the problem

- **Classification task**
- **Very small training set, called support set**
- **Goal: classifying queries**
- **Two setting:**
  - > Standard: Query seen independently
  - > Transductive: Set of queries to classify

Overview of the method

- **Training the embedding network using dense classification[1]**
- **Describing supports and query samples as collections of local features**
  - > Spatial attention mechanism
  - > Feature clustering
- **Building a graph structure with the labelled and unlabelled local features**
- **Propagating label information in the graph**

[1]Dense Classification and Implanting for Few-Shot Learning, Lifchitz et al., CVPR 2019

Local Features

- **A pixel in the feature map can be interpreted as a regional descriptor**
- **Spatial attention**
  - > Discriminative regions have high norm feature vectors
  - > A threshold on the l2-norm allow to filter out the background
- **Feature pooling**
  - > Reduces redundancy and limit the number of regions
  - > For each image, the local features are clustered using k-means clustering
  - > Local descriptors are centroid of clusters

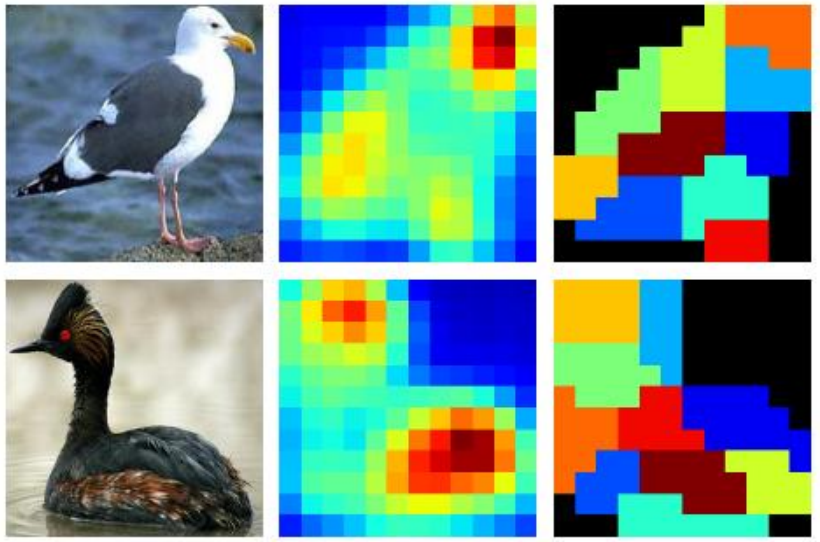

Image Norm Clusters

Propagation
Introduction

- **Graph construction**
  - > Vertices: Image samples
  - > Edge values: Cosine Similarity between samples
- **Propagation of features (optional)**
- **Propagation of label**
- **One prediction per spatial location**

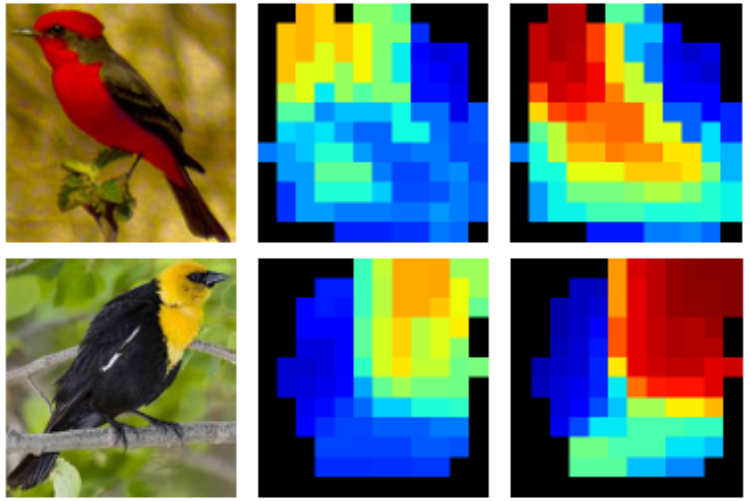

Image | Prediction Map with prototypes | Prediction map with propagation

Results

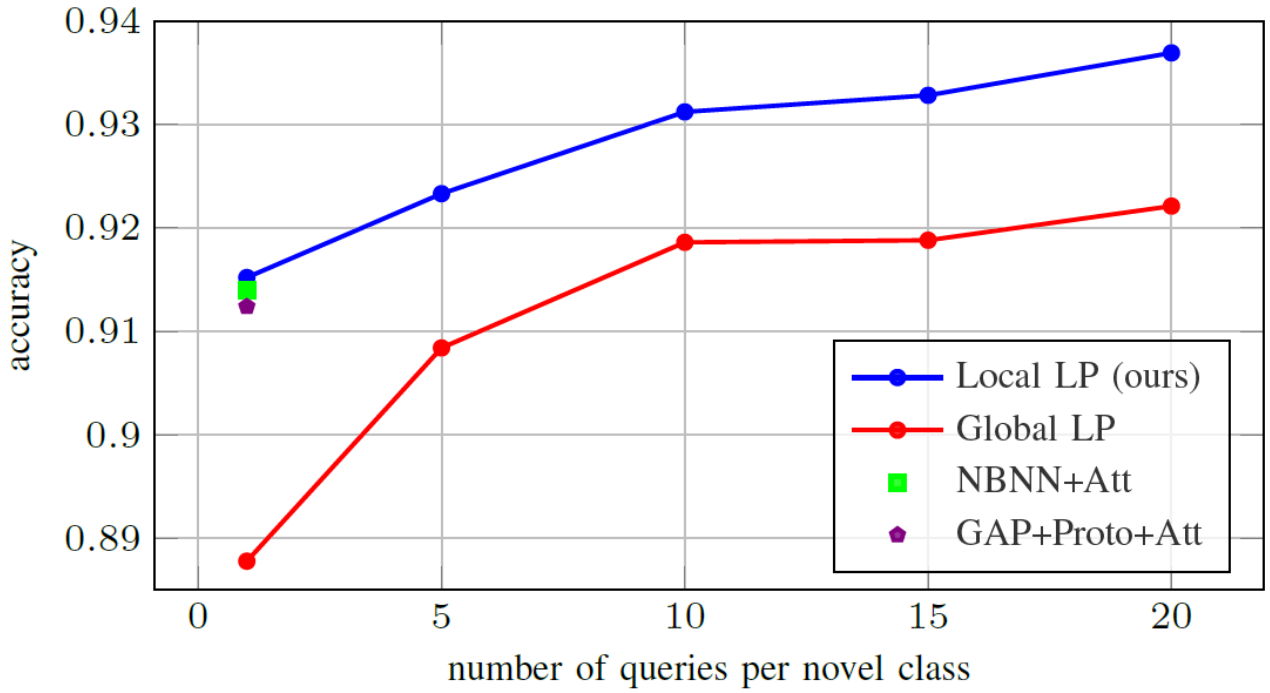

- **Applies to transductive and non-transductive setting**
- **Combined with spatial attention and spatial pooling our method is a safe choice**
  - > Performs well in the standard setting
  - > Efficiently uses unlabelled data in the transductive setting

Thank you for your attention